



OXFORD

INFINITY, CAUSATION, & PARADOX

Alexander R. Pruss

Infinity, Causation, and Paradox

Infinity, Causation, and Paradox

Alexander R. Pruss

OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,
United Kingdom

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide. Oxford is a registered trade mark of
Oxford University Press in the UK and in certain other countries

© Alexander R. Pruss 2018

The moral rights of the author have been asserted

First Edition published in 2018

Impression: 1

All rights reserved. No part of this publication may be reproduced, stored in
a retrieval system, or transmitted, in any form or by any means, without the
prior permission in writing of Oxford University Press, or as expressly permitted
by law, by licence or under terms agreed with the appropriate reprographics
rights organization. Enquiries concerning reproduction outside the scope of the
above should be sent to the Rights Department, Oxford University Press, at the
address above

You must not circulate this work in any other form
and you must impose this same condition on any acquirer

Published in the United States of America by Oxford University Press
198 Madison Avenue, New York, NY 10016, United States of America

British Library Cataloguing in Publication Data
Data available

Library of Congress Control Number: 2018939478

ISBN 978-0-19-881033-9

Printed and bound by
CPI Group (UK) Ltd, Croydon, CR0 4YY

Links to third party websites are provided by Oxford in good faith and
for information only. Oxford disclaims any responsibility for the materials
contained in any third party website referenced in this work.

Contents

<i>List of Figures</i>	xi
<i>Acknowledgments</i>	xiii
1. Infinity, Paradox, and Mathematics	1
1. Paradox and Causal Finitism	1
2. Some Mathematical and Logical Notes	4
3. Modality	7
3.1 Metaphysical possibility and necessity	7
3.2 Rearrangement principles	7
3.2.1 Defeasibility	7
3.2.2 Causal powers	9
4. Finitism: An Alternate Hypothesis	10
4.1 Time and finitism	10
4.2 Non-causal paradoxes: An advantage?	11
4.3 Mathematics: A disadvantage	13
4.3.1 Infinitely many primes	13
4.3.2 Potential infinity	15
4.3.3 *If-thenism	15
4.4 Future infinities	17
5. *Defining the Finite and the Countable	18
5.1 The finite	18
5.2 Acceptable models for the axioms of arithmetic	20
6. Evaluation	23
<i>Appendix: *Counting Future Things</i>	23
2. Infinite Regresses	25
1. How to Violate Causal Finitism	25
2. Infinite Causal Regresses	26
3. Type (i): Uncaused Regresses	27
3.1 Viciousness	27
3.2 Vicious regresses and the Hume–Edwards Principle	29
3.3 Regresses and explanatory loops	30
4. Type (ii): Causation Passing through Infinitely Many Steps	32
5. Type (iii): Outside Cause Directly Causing Each Item	33
5.1 Options	33
5.2 Regresses with outside overdetermination	35
6. *Analogy with Axiom of Regularity	36
7. Evaluation	37
<i>Appendix: *Two Kinds of Violations of Causal Finitism</i>	37

3. Supertasks and Deterministic Paradoxes	40
1. Introduction	40
2. Thomson's Lamp Revisited	40
2.1 Introduction	40
2.2 Causal finitism	40
2.3 Non-standard analysis	41
2.4 Special Relativity	42
2.5 Benacerraf's solution and the Principle of Sufficient Reason	43
2.6 Two counterfactuals	44
2.7 Evaluation	46
3. Grim Reapers	46
3.1 Introduction	46
3.2 Causal finitism	47
3.3 The absurd conclusion objection	48
3.4 A rearrangement objection	49
3.5 The mereological objection	50
3.5.1 Fusion	50
3.5.2 Necessary emergence of organic wholes	52
3.6 Uncaused lighting	52
3.6.1 Objection	52
3.6.2 The Causal Principle is true	53
3.6.3 Is the lamp lighting really uncaused?	54
3.6.4 A mysterious correlation	55
3.7 Discrete time	55
3.8 Evaluation	56
4. Infinite Newtonian Universes	56
4.1 An argument against causal finitism and a riposte	56
4.2 Smullyan's rod	58
4.3 The conditional	60
5. Another Eternal Life	60
6. Time Travel and Causal Loops	61
6.1 Grandfathers and togglers	61
6.2 Time travel and backwards causation without causal loops	63
7. Evaluation	63
4. Paradoxical Lotteries	64
1. Introduction	64
2. Countably Infinite Fair Lotteries	64
2.1 Background	64
2.2 Expected surprise	65
2.3 A guessing game	66
2.4 Symmetry	66
2.4.1 Symmetry and lotteries	66
2.4.2 *Symmetry and expected utility	68
2.5 Bayesian manipulation	71
2.5.1 The paradox	71
2.5.2 *A switchover point?	74

2.5.3 *Countable additivity and conglomerability	75
2.6 Improving everyone's chances	77
3. Constructing Paradoxical Lotteries	79
3.1 Fairness and paradoxicality	79
3.2 Lucky coin-flip sequences	79
3.3 What it is to construct a countably infinite fair lottery	81
3.4 *Coin-flips and the Axiom of Choice	83
3.5 Random walks	85
4. Objections	86
4.1 Infinite lotteries and uniform distributions	86
4.1.1 The problem	86
4.1.2 Response I: No continuous distributions	87
4.1.3 Response II: Measurement of infinite precision data	88
4.1.4 Response III: The use of the Axiom of Choice	88
4.2 *A non-normalizable quantum state	89
4.3 Limitations on our reasoning	90
5. Evaluation	91
5. Probability and Decision Theory	93
1. Introduction	93
2. Guessing with Finitely Many Errors	93
2.1 Doing a little better than one can	93
2.2 A contradiction	95
2.3 Doing much better than one can	96
2.4 *Construction of strategy guaranteeing at most finitely many errors	98
2.5 A multipersonal synchronic version	98
2.5.1 An angelic announcement	98
2.5.2 An objection and a tweak	100
2.5.3 *Making the paradox robust	101
2.6 A parody?	102
2.6.1 The story	102
2.6.2 Evaluating the parody	103
3. Satan's Apple	106
3.1 The story	106
3.2 Synchronic version	107
3.3 Diachronic version	108
3.4 Objection: Scores, desires, and promises	108
3.5 Evaluation	110
4. Beam's Paradox	111
4.1 *The mathematical formulation	111
4.2 *Synchronic version	112
4.3 *Diachronic infinite future version	113
4.4 *Diachronic supertask version	114
4.5 Evaluation of Beam's paradox	114
5. Evaluation of Decision-Theoretic Paradoxes	114
<i>Appendix: *Proof of the Theorem from Section 2.1</i>	115

6. The Axiom of Choice Machine	117
1. Less Technical Introduction	117
2. *The Axiom of Choice for Countable Collections of Reals	119
3. *Paradoxes of ACCR	120
3.1 Die-guessing games	120
3.2 Non-measurable sets	120
3.3 Banach–Tarski paradox	121
4. *An Argument for ACCR	122
5. *A Choice Machine	125
5.1 Strange mathematics and paradox	125
5.2 Coin-flips and Dutch Books	129
5.3 How to construct a Choice Machine	130
5.3.1 Angels	130
5.3.2 A four-dimensional machine	131
5.3.2.1 Making the machine	131
5.3.2.2 Using the machine	133
5.3.2.3 Causal infinitism and verifying the machine’s match	134
5.3.3 A three-dimensional machine	136
5.3.4 **Is AC needed?	136
5.3.5 Luck	137
6. Evaluation	137
Appendix: **Details of Coin-Toss Rearrangement	138
7. Refinement, Alternatives, and Extensions	140
1. Introduction	140
2. Refinement	140
2.1 Event and trope individuation	140
2.2 Histories generated by partial causal relations	142
2.3 A closer look at Grim Reapers	143
2.4 Objections to causal finitism involving partial causation	147
2.5 Absences and omissions	148
3. Some Competitors to Causal Finitism	150
3.1 Finitism	150
3.2 No infinite regresses	151
3.3 No past infinities	151
3.4 No infinite intensive magnitudes	153
3.4.1 The basic theory	153
3.4.2 Some infinite intensive magnitudes	154
3.4.2.1 Center of mass and moments of inertia	154
3.4.2.2 Mental life	154
3.4.2.3 Black holes	155
3.4.2.4 Particles	155
3.4.3 Huemer’s intensive magnitudes	156
3.4.3.1 Speed, Thomson’s Lamp, and Hilbert’s Hotel	156
3.4.3.2 *Smullyan’s rod	157
3.4.3.3 Immaterial minds	158
3.4.4 Evaluation	159
3.5 No room	159

4. Why is Causal Finitism True?	161
4.1 The question	161
4.2 Some explanatory suggestions	161
5. Further Extensions	162
5.1 Causal loops	162
5.2 Explanatory relations	164
6. Overall Evaluation	165
8. Discrete Time and Space	167
1. Introduction	167
2. Causal Finitism and Discreteness	167
2.1 The basic argument	167
2.2 From discrete time to discrete space?	167
3. Two Kinds of Discreteness	168
3.1 Subdivisibility and fixity	168
3.2 Refining the Aristotelian picture	169
3.2.1 An objection to Aristotelian discreteness	169
3.2.2 Internal and external discreteness	170
4. Physics	172
4.1 An objection to causal finitism	172
4.2 Causation and physics	172
4.3 Quantum collapse	174
4.3.1 Some background	174
4.3.2 Causation	176
4.3.3 Back to discrete time	177
5. Fields and Discrete Space	178
6. Evaluation	180
9. A First Cause	181
1. Introduction	181
2. An Uncaused Cause	181
2.1 The quick argument	181
2.2 Towards a necessary being	182
2.3 Support for the Causal Principle	183
2.4 The Kalām argument	184
3. Compatibility with Theism?	184
3.1 Theism	184
3.2 Divine motivation	184
3.3 Divine knowledge	186
3.4 Divine action	188
3.5 Limits on metaphysical possibility	191
4. Evaluation	192
10. Conclusions	193
<i>References</i>	195
<i>Index</i>	201

List of Figures

All illustrations in this volume have been compiled by the author.

1.1 Thomson's Lamp	1
1.2 Correspondence between natural numbers and even natural numbers	5
1.3 Interval notation	6
2.1 The two ways of violating causal finitism: regress (left) and infinite cooperation (right)	25
2.2 The testimony of unicorn experts	28
2.3 A theistic non-vicious regress?	33
2.4 Here, $I_1 = \{a_1, a_2, a_3\}$, $L_1 = \{e, a_1, a_2, a_3\}$, and $M_2 = \{a_4, a_5, \dots\}$	38
3.1 Some representative Grim Reaper activations	47
3.2 Some representative reversed Grim Reaper activations	48
3.3 Smullyan's rod with exponentially decreasing density and hence exponentially decreasing quasi-gravitational pull	58
4.1 A lucky case where the lottery works, with the winner being the number 2	79
4.2 A traverse of a two-dimensional array	80
6.1 The (i)–(ii) betting portfolio that you should be happy to pay a dollar for. The volume of each sphere is 1/100th of that of the cube	126
6.2 The (i')–(ii') betting portfolio that you should be happy to pay a dollar for if the argument works	127
6.3 The (i'')–(ii'') betting portfolio you should be happy to accept for free	128
6.4 A slice of a Choice Machine	132
7.1 Benardete's Boards	147
7.2 Four paradigmatic violations of (8)	163

Acknowledgments

I am especially grateful to Ian Slorach who gave me many very insightful and helpful comments and criticisms both on my blogged arguments before I started writing this book and while I was committing material to the book draft's GitHub repository. I am also particularly grateful to Miguel Berasategui, Blaise Blain, Trent Dougherty, Kenny Easwaran, Richard Gale, Alan Hájek, James Hawthorne, Robert Koons, Jonathan Kvanvig, Arthur Paul Pederson, Philip Swenson, and Josh Rasmussen. I am very grateful to other readers of my blog as well as to my audiences at Baylor University, Catholic University of America, University of Oklahoma, and the “New Theists” workshop for their patience as I tried out versions of these arguments, and for their critical commentary. Moreover, I am greatly in the debt of a number of anonymous readers of this manuscript whose careful reading has resulted in much improvement of the book. The remaining obscurities are my own accomplishment. Finally, I am very grateful to Christopher Tomaszewski for his careful work on indexing this volume.

1

Infinity, Paradox, and Mathematics

1. Paradox and Causal Finitism

A lamp is on at 10:00. Its switch is toggled infinitely often between 10:00 and 11:00, say at 10:30, 10:45, 10:52.5, and so on. No other cause affects the lamp's state besides the switch. Thus, after an odd number of toggles the light is off and after an even number it's on. What state does the lamp have at 11:00? There seems to be no answer to this question. Yet the lamp is either on or off then (Fig. 1.1).

This is known as the Thomson's Lamp paradox (Thomson 1954). Potential answers to a paradox like this fall into three general camps: logically revisionary, metaphysical, and conservative. Logically revisionary answers resolve the paradox by invoking a non-classical logic, say one in which the lamp can be both on and off at the same time, and can use the paradox as support for such revision. Metaphysical answers resolve the paradox by arguing for a substantive and general metaphysical thesis, such as that time is discrete, that there are no actual infinities, or that it is metaphysically impossible to move anything (say, a switch) at speeds whose limit is infinity (cf. Huemer 2016, 12.10.3), a thesis that explains why the story is impossible.

Conservative answers, on the other hand, refuse to revise logic or posit substantive metaphysical theses, and come in two varieties. *Particularist* conservative answers maintain that the particular story (and minor variants on it) is impossible, e.g., precisely because it is paradoxical. *Defusing* conservative answers maintain that the story as given is possible and there is no paradox in it.

A particularist answer to Thomson's Lamp paradox is simply that the story as given is impossible, since if the story were possible a contradiction would result: the lamp would be both on and off. A defusing answer, on the other hand, as given by Benacerraf (1962), notes simply that there is no contradiction in saying that the lamp is on (or off, for that matter) at 11 am: we just can't predict the state the lamp will have from the information given.

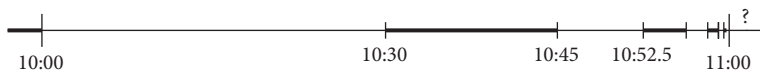


Fig. 1.1 Thomson's Lamp.

Other things being equal, conservative answers to a paradox are preferable to metaphysical ones, while metaphysical answers are preferable to logically revisionary ones. Nonetheless, other things need not be equal. For instance, while a given conservative answer may not invoke a metaphysical thesis, it may unexpectedly commit one to such a thesis, and then the benefits of conservativeness are lost. For instance, Benacerraf's solution is in tension with the Principle of Sufficient Reason. For even if there is no contradiction in the lamp's being on at 11 am, there seems to be no explanation as to why it's on then (and if it's off, there is no explanation for that).

Furthermore, if a number of paradoxes are given and each can be resolved by means of a different conservative response, nonetheless it could be preferable to resolve them all in one fell swoop by a single elegant metaphysical hypothesis that explains why *none* of the paradoxical stories are possible. For it is reasonable to prefer unified explanations of phenomena.

In this volume, I will present a number of paradoxes of infinity, some old like Thomson's Lamp and some new, and offer a unified metaphysical response to all of them by means of the hypothesis of causal finitism, which roughly says that nothing can be affected by infinitely many causes. In particular, Thomson's Lamp story is ruled out since the final state of the lamp would be affected by infinitely many switch toggles. And in addition to arguing for the hypothesis as the best unified resolution to the paradoxes I shall offer some direct arguments against infinite regresses. It is not the purpose of this book to consider *all* paradoxes of infinity—that would be an infinite task—or even all the ones that have been discovered so far. Rather, I consider a sufficient number to motivate causal finitism.¹

The availability of an elegant metaphysical solution obviates the need for resorting to logical revisionism. But we will need to be constantly on the lookout for conservative solutions to the paradoxes. Nonetheless on balance causal finitism will provide a superior resolution. Furthermore we will need to consider competing metaphysical hypotheses that resolve some or all of the paradoxes. However, it will turn out that each of the competing hypotheses suffers from one of the following shortcomings: it is broader than it should be, it fails to resolve all the paradoxes that causal finitism resolves, or it suffers from being *ad hoc*.

One can distinguish two ways of resolving a paradox: one can *solve* it by showing how an apparently incompatible set of claims is actually compatible or by showing how an apparently plausible assumption is no longer plausible after examination, or one can *kill* it by arguing that the paradoxical situation cannot occur.² In some cases, killing a paradox is not a tenable option. For instance, Zeno's paradoxes of motion can be solved, say by showing that they make assumptions about time or motion that we can reject, or they can be killed by holding that motion is impossible. Zeno, of course, wanted to kill the paradoxes, but since then most philosophers have preferred to solve them.

¹ For a more thorough survey, see Oppy (2006).

² I am grateful to an anonymous reader for this distinction.

Whether killing or solving the members of a family of paradoxes is intellectually preferable depends on the details of the situation. For instance, when the paradoxes occur in situations that we have apparent empirical observations of—arrows flying and faster runners catching up with slower ones, as in Zeno's case—killing the paradox by rejecting the actuality of the situations is apt to lead to an unacceptable skepticism, *pace* Zeno. On the other hand, when the paradoxes occur in situations which we merely intuitively think are metaphysically possible, killing the paradoxes by rejecting the metaphysical possibility of the situations may be much more tenable, since our intuitions about metaphysical possibility are unlikely to be as reliable as our empirical observations.

We may have a certain intuitive preference for solving a paradox rather than killing it. But unless the paradoxes are based on logically invalid reasoning, it will be intellectually preferable to kill all the members of a family of paradoxes in a unified way rather than solve them in a variety of different ways. One reason for this is the simple fact that to *solve* a paradox based on logically valid reasoning we have to reject a plausible premise, and hence to solve a number of such paradoxes we have to reject a number of plausible premises. But it is typically preferable to make a single assumption—especially if there is some independent reason to make the assumption beyond the need to resolve paradoxes—than to reject a number of plausible premises.

The main strategy of the book, then, will be like that of Zeno: rather than opt for a number of different solutions to different paradoxes, they will be all killed through the single assumption of causal finitism. But whereas the no-motion thesis that Zeno defends is one we have very strong empirical reasons to reject, the thesis of causal finitism is compatible with our observations (though arguing for this will take some work in interpreting modern physics).

For most of the rest of the present chapter, after some important background notes both technical and philosophical, I will consider one prominent alternate hypothesis—full finitism—and argue that in order to get out of the paradoxes, it needs to be married to a particular theory of time, the growing block theory, and that in any case it causes serious difficulties for the philosophy of mathematics. While on the subject of philosophy of mathematics, I will also offer an intriguing application of causal finitism (and of finitism as well) to the problem of defining the finite and the countable.

In Chapter 2, I will consider infinite regresses, which will give us some reason to accept causal finitism independently of the paradoxes it can kill. Then, in the succeeding chapters we will discuss several different kinds of causal paradoxes: non-probabilistic paradoxes, paradoxical lotteries, other probabilistic and decision-theoretic paradoxes, and paradoxes bound up with the Axiom of Choice from set theory. At times we will also consider what will be seen to be an analogous question: whether time travel and backwards causation are possible. I will then offer ways to refine the rough thesis of causal finitism in light of the data adduced, and argue that various alternatives to causal finitism are unsatisfactory.

Finally, I will consider two potential consequences of causal finitism. That a theory has consequences beyond what it was intended to explain gives some reason to think the theory is not *ad hoc*. At the same time, such consequences make the theory more vulnerable to refutation, since there might be arguments against the consequences.

The first apparent consequence is that time, and perhaps space as well, is discrete. If this does indeed follow, that is intrinsically interesting, but also damaging to causal finitism in that it appears to conflict with much of physics since Newton. We shall consider whether the discreteness of time actually follows and whether the kind of discreteness that is supported by causal finitism is in fact in conflict with physics, and argue that causal finitism can cohere with modern physics.

The second consequence is clearer. If causal finitism is true, then there cannot be backwards-infinite causal sequences, and hence there must be at least one uncaused cause. There is also some reason to take this uncaused cause to be a necessary being. Now the most prominent theory on which there is a causally efficacious necessary being is theism. Thus, causal finitism lends some support to theism. Interestingly, this will force us to consider whether theism doesn't in turn undercut causal finitism.

I will occasionally use the convenient phrase "causal infinitism" for the negation of causal finitism. Roughly, thus, causal infinitism holds that it is possible for something to have an infinite causal history. (Note that causal infinitism does not say that there actually *is* any infinite causal history.) Thus the point of the book is to argue for causal finitism or, equivalently, to argue against causal infinitism.

Let me end this section by noting that I do not take Thomson's Lamp to be a particularly compelling version of a paradox motivating causal finitism. There will be more discussion of it in Chapter 3, Section 2. But it is a helpful stand-in for many of the more complicated paradoxes we will consider.

2. Some Mathematical and Logical Notes

We will need some technical terminology and symbolism as general background for the book, and this will be introduced in this section. Additionally, the book contains some technical sections marked with "*" and very technical sections with "***". These can be skipped without loss of continuity. Note that any subsections of something marked with one of these markers can be presumed to have at least that level of technicality. Note also that Chapter 6 is technical or very technical as a whole apart from a less technical introduction and summary.

Start with the notion of sets as collections of abstract or concrete objects. The statement $x \in A$ means that x is a member of A . We say that a set A is a *subset* of a set B provided that every member of A is a member of B , and that A is a *proper* subset of B if it is a subset of B that does not include all the members of B . For any set B and any predicate $F(x)$ we write $\{x \in B : F(x)\}$ for the subset of B consisting of all and only the x s such that $F(x)$ (sometimes when the context makes B clear, we just write $\{x : F(x)\}$).

We can compare the cardinal sizes of sets as follows. If there is a way of assigning a different member of B to every different member of a set A (i.e., if there is a one-to-one function from A to a subset of B), then we say that $\|A\| \leq \|B\|$, i.e., the cardinality of A is less than or equal to that of B . For instance, if B is the set of real numbers between 0 and 1 inclusive, and A is the set of positive integers, then to every member n of A we can assign the member $1/n$ of B (note that if n and m are different members of A , then $1/n$ and $1/m$ are different members of B).

We say that A has fewer members than B , and we write $\|A\| < \|B\|$, provided that $\|A\| \leq \|B\|$ but not $\|B\| \leq \|A\|$. We say that sets A and B have the same cardinality when $\|A\| \leq \|B\|$ and $\|B\| \leq \|A\|$. The famous Schröder–Bernstein Theorem (Lang 2002, p. 885) says that under those conditions there is a one-to-one pairing of all the members of A with all the members of B .

Some sets are finite and some are infinite. A set is finite if it is empty or has the same size as some set of the form $\{1, \dots, n\}$ for a positive integer n . Otherwise, the set is infinite.

Infinite sets are in some ways like finite sets and in others unlike them. Both their likeness and their unlikeness to finite sets are counterintuitive to many people.

A way in which infinite sets differ from finite ones is that if A and B are finite sets with A a proper subset of B , then A always has fewer members than B . But infinite sets have proper subsets of the same size as themselves.³ For instance, if B is the set $\{0, 1, 2, \dots\}$ of natural numbers, then the proper subset $A = \{0, 2, 4, \dots\}$ of even naturals has the same size as B , as can be seen by joining them up one by one as in Fig. 1.2.

On the other hand, just as the finite sets differ among each other in size, Georg Cantor discovered that so do the infinite ones, if size is defined as above. The difference in size between infinite sets is harder to generate, however. Simply adding a new member to an infinite set doesn't make for a larger infinite set. But given a set A , we can also form the powerset $\mathcal{P}A$ of all the subsets of A . And it turns out that $\mathcal{P}A$ always has strictly more members than A —this is now known as Cantor's Theorem.⁴

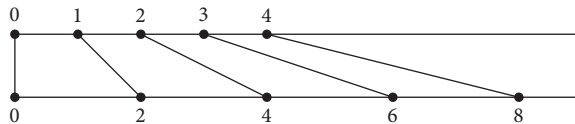


Fig. 1.2 Correspondence between natural numbers and even natural numbers.

³ **To be precise, this is only true for all Dedekind-infinite sets. If the Countable Axiom of Choice is false, then there may be infinite sets that aren't Dedekind-infinite (Jech 1973, p. 81). But standard examples of infinite sets, such as the natural or real numbers, are still Dedekind-infinite. For simplicity, I will write as if all infinite sets were Dedekind-infinite.

⁴ Here is a proof. It's clear that $\mathcal{P}A$ has at least as many members as A does, since for each member x of A , the singleton $\{x\}$ is a member of $\mathcal{P}A$. So all we need to show is that A does not have at least as many members as $\mathcal{P}A$. For a *reductio*, suppose there is a function f that assigns a different member $f(B)$ of A

If A is finite and has n members, then $\mathcal{P}A$ will have 2^n members (for we can generate all the members of $\mathcal{P}A$ by considering the 2^n possible combinations of yes/no answers to the questions “Do I include a in the subset?” as a ranges over the members of A), and $n < 2^n$. But the Cantorian claim applies also to infinite sets: in general, $\|A\| < \|\mathcal{P}A\|$.

In particular, there is no largest set. For if A were the largest set, then $\mathcal{P}A$ would be yet larger, which would be a contradiction.

Sets which are finite or the same size as the set of natural numbers $\{0, 1, 2, \dots\}$, which will be denoted \mathbb{N} , are called *countable*. An example of an uncountable set is $\mathcal{P}\mathbb{N}$. Another is the set \mathbb{R} of real numbers, which in fact has the size as $\mathcal{P}\mathbb{N}$.

A useful notation for certain sets of real numbers is given by $[a, b]$, (a, b) , $[a, b)$, and $(a, b]$ (see Fig. 1.3). Each of these denotes an interval from a to b , with the square brackets indicating that the endpoint is included in the interval and the parenthesis indicating that it's not. Thus $[a, b]$ is the set of all real numbers x such that $a \leq x \leq b$, (a, b) is the set of all reals x such that $a < x < b$, $[a, b)$ is the set of all reals x such that $a \leq x < b$ and $(a, b]$ is the set of all reals x such that $a < x \leq b$. As long as $a < b$, all the four intervals are uncountably infinite and of the same size as \mathbb{R} .

Finally, it will sometimes be convenient to talk in terms of pluralities. When I say

- (1) The members of my Department get along with each other

the grammatical subject of (1) is a plurality, the members of my Department. The verb form that agrees with that plurality, “get”, is in a plural conjugation. The subject of the sentence is not a singular object like the set of the members of my Department or some sort of a mereological sum or fusion of the members, for that would call for a singular verb, and it would make no sense to say that that singular object “get along with each other”.

We can quantify over pluralities. We can, for instance, say that for any plurality of members of my Department, the x s, there is a plurality, the y s, of people in another Department such that each of the x s is friends with at least two of the y s. Plural quantification is widely thought to avoid ontological commitment to sets. It also avoids technical difficulties with objects that do not form a set. There is no set of all sets, but it makes sense to plurally quantify and say that *all the sets* are abstract objects.

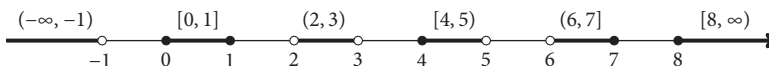


Fig. 1.3 Interval notation.

to every different member B of $\mathcal{P}A$. Let D be the set of all members x of A that are assigned by f to some set B such that x is not a member of B , i.e., $D = \{x \in A : \exists B(x = f(B) \ \& \ x \notin B)\}$. Let $x = f(D)$. Note that x is a member of D if and only if there is a subset B of A such that $f(B) = x$ and x is not a member of B . The only possible candidate for a subset B of A such that $f(B) = x$ is D , since f assigns x to D and will assign something different from x to a B different from D . Thus, x is a member of D if and only if x is not a member of D , which is a contradiction.

I will assume that pluralities have their elements rigidly. That is, if x_0 is one of the x s, then in any possible world where the x s exist, x_0 exists and is one of the x s.

3. Modality

3.1 *Metaphysical possibility and necessity*

Fairies and water made of carbon atoms have something in common: they don't exist. But there the similarity ends. For although neither is actual, the fairy is possible, while the water made of carbon atoms is not.

The kind of possibility at issue here is not merely logical. No contradiction can be proved from the existence of a fairy, but also no contradiction can be proved from the existence of carbon-based water. In each case, empirical work is needed to know the item does not exist.

There are many theories of the nature of modality.⁵ The arguments of this book are not tied to any particular such theory, but rather to intuitive judgments about cases. These intuitive judgments about cases may themselves put constraints on which theory of modality is plausible, though of course the reader will also find her theory of modality affecting what to think about the cases. That is a part of why I give so many paradoxical cases in this book: some cases may appeal to some readers while others to others.

3.2 *Rearrangement principles*

3.2.1 DEFEASIBILITY

If it is metaphysically possible to have one horse and two donkeys in a room, intuitively it's possible to have two horses and one donkey in a room as well. Lewis (1986, Section 1.8) attempted to formulate a "rearrangement principle" that justifies inferences such as this (see Koons 2014 for some more rigorous formulations). The basic idea behind rearrangement principles is that:

- (2) Given a possible world with a certain arrangement of non-overlapping spatio-temporal items, any "rearrangement" of these items that changes the quantities, positions, and orientations to some other combination of quantities, positions, and orientation that is geometrically coherent and non-overlapping is also metaphysically possible.

Unrestricted rearrangement principles sit poorly with causal finitism. Just multiply the number of button flippings and you go from an ordinary bedside lamp being turned off at night and on in the morning to Thomson's paradoxical lamp. And even if we restrict changes in quantity to be finite, an innocent forwards-infinite causal sequence can be transformed into a backwards-infinite one.

⁵ In Pruss (2011) I defend a causal powers account of modality, but nothing in the present book depends on that defense.

Nonetheless, many of our arguments for causal finitism will depend on rearrangement considerations. Isn't that cheating?

To see our way to a negative answer, observe that unrestricted rearrangement principles carry many heavy metaphysical commitments. They rule out classical theism, since on classical theism God is a necessary being, and a situation that could coexist with God could be rearranged, say by greatly multiplying evils and removing goods, into a situation that couldn't coexist with God (cf. Gulesarian 1983). They give an argument for the possibility of a universe consisting of a single walnut that comes into existence at some time, and hence for the possibility of something coming from nothing. They rule out Aristotelian theories of laws and causation on which the exercises of causal powers necessitate their effects in the absence of counteracting causes. They sit poorly with the essentiality of evolutionary origins for biological natural kinds and with the essentiality of origins for individuals. And they even rule out certain colocationist theories of material objects. For colocationists will say that wherever you have a clay statue you also have a lump of clay, but an unrestricted rearrangement principle should let you have the clay statue without the lump!

Even philosophers of a Humean bent who are uncomfortable with theism and essentiality of origins and have no problems with things coming into being *ex nihilo* need to restrict rearrangement on metaphysical grounds. For instance, Lewis (1986, p. 89) said that all rearrangements are possible "size and shape permitting". The worry is that it might turn out that material objects cannot interpenetrate, so one cannot rearrange a world with a horse standing beside a cow into a world where they occupy the same location. This seemingly purely geometrical constraint in fact needs to depend on the metaphysics of the material objects in question. Perhaps indeed a horse and a cow could not occupy the same location, but we have good reason to think that multiple bosons like photons can occupy the same location, since two bosons can have the same quantum state (Dirac 1987, p. 210). So what the "size and shape permitting" constraint comes to depends on the metaphysics of the objects, namely whether they can be colocated.

Thus, rearrangement principles like (2) should be curtailed in some way lest they ride roughshod over too much metaphysics. One way to do this is to extend Lewis's strategy by giving a list of specific metaphysical constraints like his "size and shape permitting". But it is difficult to see how we could ever be justified in thinking that our list of constraints is complete.

A better way is to stipulate that the rearrangement principles are defeasible, with the understanding that it is best if defeaters for rearrangement principles are principled rather than *ad hoc*. Theism, Aristotelian views of causation, essentiality of origins, or causal finitism could each provide principled defeaters to particular cases of rearrangement. But if one ruled out Thomson's Lamp by saying that this particular rearrangement of the ordinary bedside lamp situation is impossible, that would be *ad hoc*. If, instead, we could rule out Thomson's Lamp as well as a number of other paradoxes by means of a single general principle, namely causal finitism, that would be highly preferable. And it is this that is the strategy of the present book.

At the same time, there is always a cost to introducing another metaphysical principle like causal finitism that defeats particular instances of rearrangement. But the cost is surmountable.

I do not want the arguments of this book to be hostage to a particular rearrangement principle. Rather, I want to rely on the intuitive plausibility of the particular rearrangements that I will make use of.

3.2.2 CAUSAL POWERS

One crucial question in formulating a rearrangement principle is what properties are carried along with the objects as they are rearranged. I can rearrange a room with a braying donkey into a room with two braying donkeys. But I cannot rearrange a room with a solitary donkey into a room with two solitary donkeys. A standard thing to say is that the properties that can be carried around by the items being rearranged are the *intrinsic* properties: solitariness is not intrinsic, but braying might be. But it is notoriously difficult to define an intrinsic property (see Weatherson and Marshall 2014).

There is, however, one controversial choice that many of our arguments will require, and this is a picture of objects and their activities as having a causal nature that is carried along with their rearrangement. When one rearranges a lamp switch from one location in spacetime to another, the rearranged switch continues to have the same causal powers, and when put in the same relevant context (say, a lamp) these causal powers will have the same effects. If intrinsic properties are what can be carried along with rearrangements, then I am taking causal powers to be intrinsic properties.

This is a very intuitive picture of causal powers. It is, nonetheless, in conflict with widely held Humean views on which causal facts supervene on the global arrangement of matter in the universe. I take this conflict to provide an argument against the Humean view. The possibility of rearranging things in the world while keeping fixed the causal powers of things is intuitively more secure than the Humean theories of causation.

One strength of making a case-by-case judgment about the possibilities of rearrangements of powerful objects rather than positing a single general principle is that each such judgment can be separately evaluated by a Humean reader. The reader may decide one of three things about a particular application of rearrangement:

- (i) the application is incompatible with Humeanism and plausible enough to provide significant evidence against Humeanism; or
- (ii) the application is incompatible with Humeanism but not very plausible and so instead Humeanism provides significant evidence against this application; or
- (iii) the application can be made coherent with Humeanism, for instance, by supposing many analogous background events sufficient to ground causal laws that apply also to the rearranged case.

I leave such judgments to the reader.

4. Finitism: An Alternate Hypothesis

4.1 *Time and finitism*

Finitism holds that there can only be finitely many things (including both substances and events). Finitism, however, allows for potential infinities. Thus a collection of toy soldiers to which a new toy soldier will be added every day would be potentially infinite as for any number n , it would eventually have more than n elements. But according to finitism there are no actual infinities. There are always only finitely many things.

Finitism has an impressive philosophical history, going back at least to Aristotle's responses to Zeno's paradoxes,⁶ and being the generally accepted philosophical orthodoxy in the Middle Ages.

The exact upshot of finitism depends on which theory of time it is combined with.

The *eternalist* thinks of past, present, and future things as all ontologically on par, and believes that (barring some catastrophe) our great-great-great-grandchildren exist and Alexander's great warhorse Bucephalus also exists. Of course, the great-great-great-grandchildren and Bucephalus don't *presently* exist. But they nonetheless really do exist. The *growing block theorist* takes reality not to extend to the future, but to include the past and present.⁷ Thus our great-great-great-grandchildren don't exist (though it might be true that they will exist), but Bucephalus does. The *presentist*, on the other hand, only accepts presently existing entities as existing. I will, further, take all of these theses about time to claim to be necessarily true.

Finitism plus eternalism straightforwardly entails causal finitism: if there can only be finitely many things, and that includes past, present, and future, then of course nothing can be affected by infinitely many causes. Thus any paradox ruled out by causal finitism will be ruled out by finitism plus eternalism. But unfortunately finitism plus eternalism also entails that the future must be finite—that there cannot be infinitely many future events. But surely it is possible to have an infinite future full of different events or substances, say with a new toy soldier being produced every day forever. Thus, finitism is implausible given eternalism.

Given presentism, on the other hand, finitism is compatible with infinite sequences of causes, as long as at no particular time are there infinitely many causes. Thus, finitism plus presentism does nothing to rule out the infinitely many toggles of the switch in Thomson's Lamp. While the other paradoxes have not yet been discussed, many of them will also have the diachronic character of Thomson's Lamp and hence will be untouched by presentist finitism.

⁶ Of course, in a degenerate way, Parmenides was a finitist, as he thought that there could be only one thing.

⁷ There is also a variant by Diekmeyer (2014) that includes only the past. That variant will not be helpful to the finitist, and I will stick to the canonical version that includes the present.

This leaves growing block plus finitism. If it's necessarily true that a cause is earlier than or simultaneous with its effect, then growing block finitism does entail causal finitism, and hence can rule out all the paradoxes that causal finitism can. Given the combination of (a) growing block theory, (b) finitism, and (c) the thesis that causes are either temporally prior to or simultaneous with their effects, we do get causal finitism again, and hence we can rule out all the paradoxes that causal finitism can rule out.

The finitist's best bet at paradox removal is thus to adopt growing block together with the thesis that causes are prior to or simultaneous with their effects.

Unfortunately, there is a powerful argument against growing block theory due to Merricks (2006). Many people have thought thoughts about what date or time it is, thoughts expressible in sentences like: "It is now 2012" or "It is now noon." If growing block theory is true, many of these thoughts are in the past, and most have a content that is objectively false. On the growing block theory, the "now" is the leading edge of reality, the boundary between the real and the unreal. The thought expressed by "It is now 2012" is true if and only if 2012 is at the leading edge of reality. But 2012 is not at the leading edge of reality. Moreover, my present thought that it is now 2018 has no better evidence than the "It is now 2012" thought that was in 2012. Since most thoughts of this sort, with the usual sort of evidence for them, are false, I should be skeptical about whether it is now 2018. And that's absurd. Presentism escapes this argument by denying that the past thoughts exist. Moving spotlight versions of eternalism, on which there is something like an objective "moving spotlight" illuminating the "now" are also subject to this objection: most of the "It is now t " thoughts are not illuminated by the spotlight, and yet "now" implies such "illumination". But B-theoretic eternalism (e.g., Mellor 1998), which claims that the "now" is a mere indexical, rather than an expression of an objective changing property (like being at the leading edge of reality or being "illuminated"), is not subject to the objection. Hence, the finitist's best bet at paradox removal requires adopting a particularly vulnerable theory of time.

We now consider more advantages and disadvantages of finitism *vis-à-vis* causal finitism as a way out of paradoxes.

4.2 *Non-causal paradoxes: An advantage?*

Imagine Hilbert's Hotel—a hotel with infinitely many rooms numbered 1, 2, 3, You can have lots of fun with that. Put a person in every room, and then hang up the sign: "No vacancy. Always room for more."⁸ When a new customer asks for a room, just put them in room 1, and tell them to tell the person in the room to move to the next room, and to pass on the same request. You can even have infinitely many people vacate the hotel and still have it full. If all the people in the odd-numbered rooms leave, you can tell each person in the even-numbered rooms to move to a room whose number is half of their room number.

⁸ The sign suggestion comes from Richard Gale.

While the stories I gave involved causation, that was only for vividness. To see the paradoxicality, all we need to notice is that the guests can move around (even causelessly, if that's possible) to make space for a new guest, and that every second guest can leave, while the hotel quickly regains its fullness.

The root of these paradoxes is that an infinite collection can be put in one-to-one correspondence with a proper subset (cf. Fig. 1.2, above). If finitism is true, then of course we can rule all such paradoxes out of court. This provides a simple argument for finitism: If finitism is not true, then Hilbert's Hotel is possible. But Hilbert's Hotel is absurd and hence not possible, so finitism is true.

But while Hilbert's Hotel is indubitably *strange*, the strange and the absurd (or impossible) are different, as is proved by the strangeness of the platypus. We could just conclude from Hilbert's Hotel that infinity is roomier than we previously thought.

One might think that an outright contradiction can be proved from Hilbert's Hotel. For instance:

- (3) The collection of even-numbered rooms is smaller than the collection of all rooms.
- (4) The two collections can be put in one-to-one correspondence (matching room n for even values of n with room $n/2$).
- (5) Two collections that can be put in one-to-one correspondence are of the same size.
- (6) If A is smaller than B , then A is not the same size as B .
- (7) So the two collections both are and are not the same size.

Assuming that "smaller" and "same size" are used univocally throughout, there are two ways of rejecting the argument. First, one can reject (3). Granted, for *finite* sets a proper subset is smaller than its proper superset. But we shouldn't expect this to be true in infinite cases. After all, infinite cases are different from finite ones. Alternately, one can reject (5) (of course, here, "size" cannot be stipulated as was done in Section 2).

The question of how to gauge what is absurd and what is merely strange is a difficult one. While I do not see much of a cost to rejecting (3) or (5), others will. Nonetheless, in the case at hand there is a very strong reason to reject finitism, and thus accept the possibility of something like Hilbert's Hotel.

A related paradox is the following. Intuitively, there are more positive integers than prime numbers. But now imagine an infinite collection of sheets of paper, with one side red and one side green. It is as clear as anything that the number of red sides equals the number of green sides. Now suppose that the green sides are numbered⁹ 1, 2, 3, ..., and suppose that the red side of a piece of paper that has n on its green side contains an inscription of the n th prime number (we won't run out of primes:

⁹ The "ink" will have to be a non-molecular sort, since in order to fit very long numbers on the page, the numerals will have to get smaller and smaller.

there are infinitely many of them¹⁰). Then the number of red sides equals the number of primes, and the number of green sides equals the number of positive integers, and since the number of red sides equals the number of green sides, we conclude that the number of primes equals the number of positive integers, which contradicts the assumption that there are more positive integers than primes.

Denying the possibility of an actual infinite kills the paradox. But one can also solve the paradox by saying that the argument is a *reductio ad absurdum* of the initial intuition that there are more positive integers than prime numbers. And denying the possibility of an actual infinity only kills the paradox at the cost of undercutting this initial intuition in a different way. For if actual infinities are impossible, then it seems to make no sense to say that there are more positive integers than prime numbers, since neither infinite plurality can actually exist. We will now consider an argument *against* finitism along a similar line.

4.3 Mathematics: A disadvantage

4.3.1 INFINITELY MANY PRIMES

A flatfooted mathematics-based argument against finitism is: “There are infinitely many primes. So finitism is false.” (Note, too, how this argument doesn’t seem to affect *causal finitism*, since numbers appear to be causally inert.)

But perhaps when the finitist told us that there couldn’t be an actual infinity of things, she was thinking about *concrete* things like rooms and not abstract ones like numbers?

This is not plausible, however. For while the finitist’s arguments are formulated in terms of concrete things, the intuitions on size that underwrite arguments like (3)–(7) apply just as much in the case of abstracta. The fact that there are as many even numbers as natural numbers is in itself counterintuitive, and the Hotel merely makes this more vivid. Thus the finitist cannot afford to restrict her view to concrete entities, since doing so leaves unanswered paradoxes that intuitively are of exactly the same sort as the ones she resolves.

Furthermore, the distinction between abstract and concrete objects is not a particularly clear one, and unless it is clarified it is hard to say why exactly one would think that infinite collections of concrete things are a problem but infinite collections of abstract things are not.

For instance, one clarification of the notion of concreteness in the literature is given by Pruss and Rasmussen (2018), who say that an entity is concrete if and only if it is possible for the entity to cause something. If this is right, then the finitist who

¹⁰ The classic proof is a *reductio ad absurdum*. If there are finitely many primes, let p be the product of all of them. Then $p + 1$ is bigger than every prime. Moreover, $p + 1$ is not divisible by any prime, since it yields remainder one when divided by any prime. But a number not divisible by any prime is a prime. So $p + 1$ is a prime bigger than every prime, which is absurd.

thinks that there are infinitely many primes would need to answer why it is that the possession of causal powers rules out infinities. After all, although the rooms in Hilbert's Hotel presumably have the possibility of causing things (for instance, a wall can cause pain in a fist), nothing in the paradox depended on the causal possibilities for the rooms. One might as well say that Hilbert's Hotel is impossible because the rooms possibly have color.

So objecting to the flatfooted argument from primes on the basis of a distinction between abstracta and concreta is not promising. A better option, however, is to point out that the argument depends on a Platonist interpretation of the sentence "There are infinitely many primes." But Platonism is not the only position in the philosophy of mathematics. There are other options.

But not all of the other options are available to the finitist. The finitist labors under the special disability that not only does she think that the infinitude of mathematical objects does not exist, but she thinks that nothing with the relevant structure—a central aspect of that structure being infinitude—could possibly exist. Thus the mathematician is someone studying impossible situations. But while it has always been a wonder that something as rarefied and abstract as mathematics should be applicable to the actual world, it is a true miracle that the study of genuinely impossible things should be of such relevance to us. How could the queen of the natural sciences be the study of impossible structures?

Consider, too, that mathematics involves proofs from axioms. Certain axioms are controversial, and hence they are not always assumed. Thus, in some contexts, mathematicians go out of their way to flag that they did not assume the Axiom of Choice in a proof. The reason for excluding axioms from the assumptions for a proof is two-fold. The first reason comes from epistemic concerns about the truth of the axiom. Since we're not sure the Axiom of Choice is true, it is safer not to assume it. The second is theoretical: even if the given axiom were true, we would be interested in knowing what it would be like if we had a system that did not satisfy that axiom.

But the finitist takes herself to have established on pain of absurdity that there cannot be infinitely many things. If that's really established, then we do not have an epistemic reason to exclude the Axiom of Finitude—that there are only finitely many things—from the axioms used in our mathematical work. And if the Axiom of Finitude is necessarily true, then the question of what it would be like if a system failed to satisfy the axiom is more of a logician's or philosopher's question than a question for the typical mathematician to wonder about. While some mathematicians do in fact study alternative collections of axioms for set theory, typical mathematicians are happy to assume axioms of set theory they find intuitive. Likewise, if the Axiom of Finitude were necessarily true, then most working mathematicians should assume it and study its consequences, rather than engaging in the study of the *per impossibile* counterfactual of what would happen if there were infinities. The resulting mathematical practice would be very different from ours, and we have good reason to doubt the

wisdom of this bold proposal given the fruits that current mathematical practice has brought.

4.3.2 POTENTIAL INFINITY

The infinitude of the primes, while striking, is perhaps not the biggest difficulty for the finitist. A finitist can adopt the view that the natural numbers, and infinite subsets of it like the primes, can be taken to be potentially infinite, not in the sense that there is a potential for generating the whole set, but rather in the sense that no matter how many members of the set have been generated, another can be added (see Craig 1979, p. 200 and Oppy 2006, pp. 263–4). This could perhaps even be done in a nominalist-friendly way: No matter how many boxes one has with a prime number of oranges in each, one could create another box with a different prime number of oranges.¹¹

But we have a more serious problem. For mathematics doesn't limit itself to claims about *particular* infinite sets of natural numbers. It also talks about sets of sets of natural numbers. Consider for instance the powerset $\mathcal{P}\mathbb{N}$ of the set of natural numbers \mathbb{N} . This is the set of all sets of natural numbers (i.e., the set of all sets whose members are natural numbers). This powerset is infinite, but Cantor's Theorem tells us that it is an uncountable set, and hence cannot be generated by a sequence of successive addition of the sort that could be used to generate the set of all prime numbers. Nor are sets like that at all rare in mathematics: everyday mathematical practice engages in abstraction upon abstraction, happily studying sets of sets of sets at multiple levels.

One could try to identify sets of natural numbers with methods for generating a set by successive addition.¹² But since there are uncountably infinitely many sets of natural numbers, there will have to be uncountably infinitely many such methods. And that would directly violate finitism. And for the Cantorian reason one can't just talk of a method of generating methods: there are too many methods for the set of methods to be generated simply by adding one more at a time, in the way that the set of prime numbers of boxes of oranges can be.

4.3.3 *IF-THENISM

Obviously, finitism undercuts mathematical Platonism. For a concrete example of how finitism can undercut another plausible view in the philosophy of mathematics, let's consider if-thenist philosophy of mathematics which holds that the discoveries of mathematics are necessary conditionals such as: "Necessarily *if* these axioms hold of a system, *then* in that system there are infinitely many entities that count as primes."

¹¹ This formulation not only is compatible with finitism but presupposes it, by assuming that there are only finitely many boxes. For if we had an infinite number of such boxes, it could be the case that all prime numbers were represented among the box–orange counts.

¹² I am grateful to Blaise Blain for this suggestion.

No claim is made that these axioms hold of a system and hence no claim is made that there are actually infinitely many things.

But a proposition p entails a proposition q provided that in every possible world where p is true, q is true as well, or, equivalently, provided that there is no world where p is true but q is not. Thus an impossible proposition entails every proposition, since there is no world where an impossible proposition is true. But the finitist claims it is impossible for there to be an infinite number of things. Thus any axioms that entail that there are infinitely many primes entail every proposition. Hence, given finitism, not only is it true that necessarily if the axioms of arithmetic are true then there are infinitely many primes, but it's also true that necessarily if the axioms of arithmetic are true then an impossibility happens, and thus that necessarily if the axioms of arithmetic are true, then there are square circles and round triangles.

The finitist if-thenist might want to distinguish between necessities. Perhaps there is metaphysical necessity—the necessity in saying that water is H_2O and that nothing can be its own cause—and narrowly logical necessity, namely provability in a logical system. Mathematicians are not interested in the claim that it's *metaphysically* necessary that if the axioms hold, then there are infinitely many primes. Rather, their interest is in the claim that the conditional is *logically* necessary, that one can *prove* the existence of infinitely many primes from these axioms.

This pushes the if-thenist into a version of logicism that holds that the discoveries of mathematics are about logical provability. Logicism is widely held to have been discredited by Gödel's incompleteness theorems (Hellman 1981). But the version of logicism that has been directly discredited is one that identifies mathematical truth with provability, given that Gödel showed that in every (recursive) logical system that contains enough arithmetic there will be unprovable truths.

The version of logicism under discussion is different. No claim is made that mathematical truths are provable propositions. Rather, the present version of logicism holds that what mathematicians discover is simply that some sentences (say, certain material conditional sentences) can be proved. The real subject matter of mathematics on this view is proofs or provability rather than mathematical truth.

But this is seriously implausible. Typically we are interested in proofs not for the sake of the meta-mathematical fact that something can be proved, but for the sake of knowing that the thing that can be proved—perhaps a conditional, if if-thenism is right—is in fact true. It is surely the truth of mathematical conditionals that makes them applicable, and proof is but a tool for learning about truth.

Moreover, proof is not the only technique used by mathematicians. Many engage in numerical experimentation. For instance, Goldbach's Conjecture says that every even integer greater than two is the sum of two primes. This has been verified by computer, that every even integer greater than two up to 4×10^{18} is the sum of two primes (Silva 2015). There is a straightforward inductive argument (in the science, not mathematics, sense of “inductive”) from this that all even integers greater than two are like that. However, there is no corresponding inductive argument from the computer

experiments to its being *provable* that all even integers have that property.¹³ Indeed, the computer scientist Donald Knuth took seriously the possibility that Goldbach's Conjecture is true but not provable.¹⁴ We can best make sense of numerical experimentation if we take mathematics to be a search for truth rather than provability.

In other words, in the spirit of if-thenism we can suppose mathematics is about entailment or about provability. The former explodes (all truths are entailed by axioms that imply infinities) if the finitist is right that actual infinities are impossible, while the latter is unsatisfactory in general.

4.4 Future infinities

In order to have a theory that is at all plausible, the finitist has to allow for the possibility of future infinities. It is surely possible that a coin will be flipped infinitely often, say, or that someone will live an infinite number of days (say, in an afterlife). In order for finitism to allow for the metaphysical possibility of such scenarios, we had to couple finitism with a growing block theory of time (or presentism—but, as noted, that doesn't help with enough paradoxes) rather than eternalism.

But at least some of the counting paradoxes motivating finitism can be run in a diachronic futureward way. Take first the paradox about natural numbers and primes that I made vivid with colored pieces of paper. Suppose now that every day for an infinite amount of time a new piece of paper will be produced, green on one side and red on the other. Suppose, further, that the green side of the n th piece of paper will have n written on it, and the red side will have the n th prime written on it.

It surely makes sense to talk of how many events of some type *will* happen over an infinite future.¹⁵ We can now run this argument:

- (8) The number of green sides produced will equal the number of red sides produced.
- (9) The number of green sides produced will equal the number of positive integers written down.
- (10) The number of red sides produced will equal the number of primes written down.
- (11) All the positive integers and all the primes will be written down.
- (12) So, the number of positive integers equals the number of primes.

¹³ For any n for which Goldbach's Conjecture is true, it's provable that Goldbach's Conjecture is true for n —all it takes to prove it is a lot of long multiplication to check that two prime summands that n decomposes into are indeed primes. Hence, there is an inductive argument that for each n , it is provable that Goldbach's conjecture holds for n . But it does not seem to follow from the claim that Goldbach's Conjecture is provable for each particular n that it is provable that it holds for all n .

¹⁴ See Knuth (2002). A similar issue arises with the Riemann Zeta Hypothesis (RZH), and the logician Martin Davis has speculated that RZH is a good candidate to be unprovable (Jackson 2008, p. 571).

¹⁵ Especially if the events are produced by some deterministic process, as in this case we can imagine to be the case.

Perhaps a growing block theorist will deny that statements about counting future events or objects, at least when that number is infinite, make any sense. But that would be a *reductio ad absurdum* of the position. For surely it makes perfect sense to say that if I were to toss a fair coin infinitely many times, it would likely come up heads infinitely often. In any case, in fact, we *can* make sense of counting future things given either presentism or growing block. The technical details are given in the Appendix to this chapter. Thus, future things can be counted even if we do not suppose their actual existence. Paradoxes of infinite counting, thus, are only eliminated by supposing finitism when that finitism also requires finite numbers of future objects—and that’s not a plausible finitism.

5. *Defining the Finite and the Countable

5.1 *The finite*

The standard mathematical definition of a finite set is that S is finite if and only if S has the same cardinality as the set $\{0, 1, 2, \dots, n\}$ for some natural number n . One difficulty with this definition is that it presupposes the concept of a natural number, and one may reasonably worry that natural numbers are precisely the numbers that count the elements in a *finite* set.

Granted, we can state explicit axioms that natural numbers satisfy, say the Peano Axioms. But the natural numbers are not the only mathematical structure that satisfies these axioms. For instance, hypernatural numbers satisfy the same axioms (cf. Robinson 1996), and yet include numbers that from our intuitive point of view are infinite, and by Gödel’s First Incompleteness Theorem, any similar set of axioms will fail to fully characterize the natural numbers.¹⁶ Another difficulty is that the notion of sameness of cardinality depends on one-to-one correspondences, which depends on the background set theory—and again we have the problem that any background set theory can be extended.

Alternately, we might try to characterize a finite set as one that gets smaller when an item is removed, i.e., a set whose elements cannot be put into a one-to-one correspondence with those of a proper subset.¹⁷ This definition, too, depends on the background set theory for the notion of “one-to-one correspondence”.

One may be satisfied with the above definitions. But causal finitism makes possible a metaphysical rather than set-theoretic definition. Given causal finitism any plurality of objects that has a joint effect must be finite. Then note that if the y s are a finite plurality of objects, and there is a relation R such that for each x_0 among the x s there is a y_0 among the y s such that x_0 stands in R to y_0 and nothing else among the x s stands in R to y_0 , then the x s are a finite plurality as well.

¹⁶ Second-order arithmetic does characterize the natural numbers, but only relative to a background set theory (cf. Shapiro 1985, p. 735).

¹⁷ As noted earlier, this only defines a Dedekind-finite set, but given a weak version of the Axiom of Choice, it is equivalent to the usual natural number-based definition.

We can now offer this definition:

- (13) There are finitely many xs if and only if possibly there are ys such that
 (a) possibly the ys are all in the causal history of a single item and (b) possibly
 there is a relation R such that for each x_0 among the xs there is a y_0 among
 the ys such that x_0 stands in R to y_0 and nothing else among the xs stands
 in R to y_0 .

Assuming causal finitism, we can argue that if the xs satisfy the right-hand side of (13), there are finitely many of them. Condition (a) and causal finitism guarantees that there are finitely many ys in some possible worlds. But since the elements of a plurality do not vary between possible worlds where that plurality exists, and since finiteness intuitively depends only on the elements of a plurality, it follows that there are finitely many ys in every possible world where the ys exist. Condition (b) then guarantees that in some possible world there are no more xs than ys , and hence that there are only finitely many xs there. And by the same reasoning as before, it follows that there are actually only finitely many xs .

The converse depends on some plausible theses about what is possible. Plausibly, if there are finitely many xs , it is possible to have a finite number of minded beings, the ys , such that each of the ys thinks about exactly one of the xs and each of the xs is thought about by exactly one of the ys . It is also plausible that we can further require that the minded beings produce a single effect together—maybe a committee statement about the xs . Then if we let R be the “is thought about” relation, the right-hand side of (13) will be satisfied.

Thus, given causal finitism, (13) is true, and as causal finitism is a claim about what is necessarily the case, the argument can be run in any world, and so (13) is necessarily true. Thus, given causal finitism, (13) is intensionally correct. Whether it is a good definition is another question.

And as a bonus once we have a metaphysical definition of the finite, we can get a metaphysical definition of the countable. A total ordering relation is a transitive (if $a \leq b$ and $b \leq c$ then $a \leq c$), antisymmetric (if $a \leq b$ and $b \leq a$ then $a = b$), and total ($a \leq b$ or $b \leq a$ for all a and b) relation. Then:

- (14) There are countably many xs if and only if it is possible to have a total ordering relation \leq on the xs such that for any b among the xs there are only finitely many a among the xs such that $a < b$.

Note that (13) seems to trivialize causal finitism. For given (13), it is trivial that if the ys are in the causal history of a single item, then there are finitely many of them—just let the xs be the ys and let R be identity. But the trivialization is not complete. For while the statement of causal finitism is trivially true, what is not trivially true is that the definition of the finite satisfies our intuitive beliefs about the finite.

One can do a similar thing with finitism in place of causal finitism. In fact, given finitism, we can remove condition (a) in (13). But we have seen that finitism is not satisfactory.

5.2 *Acceptable models for the axioms of arithmetic*

One of the deepest problems in the philosophy of mathematics is how our minds connect with the mathematical realities. What makes our word “seven” connect with one particular mathematical object and our phrase “the set of primes” connect with another particular mathematical object? One version of the difficulty is Benacerraf’s (1965) famous identification problem. There are infinitely many sets of mathematical entities that could play the role of the set of natural numbers such that the very same mathematical truths will hold of all of them. For instance, we could identify the natural numbers $0, 1, 2, 3, \dots$ with the sequence $\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}, \{\emptyset, \{\emptyset, \{\emptyset\}\}\}, \dots$, where the first item in the sequence is the empty set and each item in the sequence is the set of all the preceding items, and then define the arithmetical operations accordingly. Alternately, we can identify the natural numbers with the sequence $\emptyset, \{\emptyset\}, \{\{\emptyset\}\}, \{\{\{\emptyset\}\}\}, \dots$, and define arithmetical operations accordingly. On both definitions, there will be infinitely many primes, Fermat’s Last Theorem will hold, and generally speaking all arithmetical truths true in one of these settings will be true in the other, because there will be an isomorphism (technically, elementary equivalence) between them.

But while Benacerraf’s identification problem is puzzling for the philosophy of mathematics, it is not much of a problem for the typical working mathematician. It does not matter to the typical working mathematician which construction of the natural numbers is used, because they all give rise to the same arithmetical truths. The working number theorist can say that she is happy to have her number theory interpreted in any of these set-theoretic models of the natural numbers. The models are all equivalent, so it does not matter which one we choose.

There is, however, an identification problem that cuts more deeply than Benacerraf’s. From Gödel’s First Incompleteness Theorem (Boolos and Jeffrey 1995, p. 188) together with the Soundness Theorem, we learn that consistent (recursively specifiable) axioms sufficient for doing arithmetic underdetermine arithmetical truths, so that for any such axioms there will be a sentence s that is true in some models of the axioms and false in others. It is not true that all the relevant set-theoretic models are isomorphic, and so we cannot blithely say that it does not matter which one we choose. There are infinitely many models that are mutually non-isomorphic.¹⁸

There still is a retreat available to the working mathematician. She can say that she is confident of her favorite axioms, say the Peano Axioms of number theory or the Zermelo–Fraenkel Axioms of set theory, and all she is interested in is the question of what can be proved from these axioms.

But incompleteness does not let one off the hook so easily. For *provability* is itself a mathematical notion, and Gödel’s proofs of his incompleteness theorems proceeded by showing that one can encode sentences and proofs as natural numbers.

¹⁸ This follows by iterating applications of First Incompleteness and Soundness.

A simple example is to use an alphabet and symbol set of 90 characters or less, and then encode each letter and symbol as a two-digit sequence between 10 and 99 ($A = 10$, $B = 11$, $C = 13$, etc.), and encode an arbitrary finite sequence of letters and symbols as a decimal number consisting of strings of these two-digit sequences (e.g., $ABCA = 10111310$). The property of being a well-formed formula or sentence becomes equivalent to a property of the encoding number that can be put in the language of arithmetic, and the property of being a proof of some sentence becomes equivalent to an arithmetical relationship between the number encoding the alleged proof and the number encoding the concluding sentence.

A lesson from the encoding is that the question of what can be proved from some (recursively specifiable) axioms itself is an arithmetical question. Thus, if we have an underdetermination of the model of arithmetic, we run the danger of having a corresponding underdetermination of what can be proved from what, and the retreat to provability does not help.

One might hope that the danger does not eventuate, because perhaps all the models of arithmetic that agree on our favorite axioms—say, the Peano Axioms—will also agree on what can be proved from what. However, Gödel's Second Incompleteness Theorem punctures that hope. For according to the Second Incompleteness Theorem, a consistent (recursively specifiable) axiomatization sufficient for doing arithmetic cannot prove that it is consistent, i.e., it cannot prove an arithmetical statement that is equivalent to the statement that the axiomatization is consistent. It follows from Second Incompleteness that there is a model of the axioms according to which the axioms are not consistent. But if the axioms are in fact consistent, there is also a model—intuitively, a *correct* model—according to which the axioms are consistent. In other words, models of the natural numbers disagree on what axiomatizations are or are not consistent. But the question of which axiomatization is or is not consistent is precisely a question about what can be proved from what. For an axiomatization is consistent if and only if an explicit contradiction (say, that $0 = 0$ and $0 \neq 0$) cannot be proved from it. Hence, assuming our axiomatization is consistent, its models disagree on whether a contradiction can be proved from it, and in particular disagree on provability.

Nor can we say that what the models disagree on is mathematically or philosophically uninteresting. For as we saw, what the models disagree on is one of the most interesting questions in mathematics and the philosophy thereof: whether our favorite axioms of arithmetic are consistent. If we do not have a way of connecting our thoughts about numbers (or, equivalently, sentences) to a family of models of arithmetic that agree on the question of consistency, we have to say that the question of consistency is literally nonsense—and that does not seem plausible.¹⁹

¹⁹ One could, of course, add to the Peano Axioms the statement that the Peano Axioms are consistent. But then the same Second Incompleteness argument will apply to the Peano Axioms supplemented with the claim that they are consistent: the supplemented axioms will underdetermine whether the supplemented axioms are consistent, assuming they in fact are.

The retreat to what can be proved from the axioms thus does not escape the problem of disagreeing models of arithmetic. We might be able to solve the problem by saying that our interest in number theory is neither in discovering what is *true* of the natural numbers nor in discovering what is *provable* about them, but simply in actually generating proofs. This seems to be a costly retreat that undercuts the motivations of mathematical curiosity.

The causal finitist account of the finite from Section 5.1, however, goes some way towards solving the underdetermination problem for arithmetic. To motivate this, start with the intuitive thought that there is such a thing as “genuine” natural numbers, and hence that there is such a thing as “genuine” provability, i.e., provability by means of proofs that can be encoded as genuine natural numbers. But there are models of arithmetic called “non-standard”, which add to arithmetic natural numbers that from our “genuine” point of view are infinite, but nonetheless satisfy the Peano Axioms.²⁰ These infinite natural numbers can then be thought of as encodings of “infinite proofs”.²¹ And it is no surprise that provability depends on what, if any, infinite proofs are allowed.²² This intuitive thought suggests that to get out of the problem, we need a way of limiting our numbers and proofs to be finite.

And we saw in Section 5.1 that the causal finitist (or the finitist *simpliciter*) can give an account of the finite. We can now specify that we only count as acceptable those models M of arithmetic that have this property:

- (15) For any natural number n of M , there are only finitely many natural numbers of M less than n , where “finitely many” is understood metaphysically according to the account of Section 5.1.

This does not solve the Benacerraf Identification Problem: if one model satisfies (15), so will any model isomorphic to it. It does not determine all the axioms of set theory, since one can have inequivalent models of set theory that share the same natural numbers. But what it can help with is the most radical aspect of the underdetermination problem, the problem of models that disagree on the natural numbers and hence on provability. For it is a quite plausible metaphysical thesis, a thesis apparently unaffected by the Incompleteness Theorems, that any two models of Peano Arithmetic that satisfy (15) are isomorphic, and hence at least agree on provability.

Of course, we can plug *any* metaphysical account of the finite into (15). But the causal finitist account just given is conveniently available.

²⁰ The non-standard arithmetic of Robinson (1996) will probably be most familiar to philosophers, but does not exemplify the relevant variation with respect to provability.

²¹ Cf. the phrase “non-standard proof” in Kunen (1980, p. 146).

²² Famously, Leibniz thought that there were infinite proofs, and indeed he based an account of the difference between contingency and necessity on the difference between what has infinite and what has finite proof. However, Leibniz thought that there was an absolute notion of infinite proof, whereas the lesson from Second Incompleteness is that what infinite proofs there are will vary between models.

6. Evaluation

Given some additional assumptions, finitism implies causal finitism, and hence has all of the paradox-excluding benefits of finitism. But finitism is an overreach. The absurdities claimed as supports for finitism can be reasonably taken to simply be indicators of the strangeness of infinity. Causal finitism is a more modest thesis and that is a reason to prefer it.

Furthermore, finitism endangers mathematics. Not only are widely studied mathematical structures like the natural numbers not realized, but given finitism they are impossible things, like a water molecule with three hydrogen atoms or a self-caused rock. Causal finitism has no such problems. It is compatible with an infinitude of causally inert Platonic entities, and it is compatible with the existence of an infinitude of concrete entities that are, say due to their spatiotemporal arrangement, unable to cooperate causally together.

Next, if one has reason to accept eternalism—the full reality of both past and future events and objects—then one must reject finitism. For it is metaphysically possible to have a universe that will go on for an infinite future to spawn new entities, and given eternalism this possibility is incompatible with finitism.

In this chapter, we also saw an interesting application of causal finitism: it allows for a metaphysical definition of the finite. That application shows that there is some theoretical benefit to causal finitism for the philosophy of mathematics, and hence provides some evidence for causal finitism. In the rest of the book, I hope to increase the level of evidence to one that I take to be quite compelling, primarily by solving a number of interesting paradoxes.

Appendix: *Counting Future Things

We can give a precise definition of what it means to say that there will be n things (objects, events, etc.) of some kind K , even for n infinite, without any commitment to the actual existence of future things. For simplicity, I will assume that K is a kind such that if a thing is ever a K it is always a K . (If we want to count members of some stage-kind, like students, we can simply replace a stage-kind K_0 with the non-stage-kind K_0 -at-some-future-time-or-other.) The definition I will give will work both given presentism or growing block.

To proceed, let F be the set of all sets of times containing at least one future time. Suppose $S \in F$. Then S is a set of times with at least one future time being a member of S . Next, say that a presently existing x *exactly temporally occupies* S provided that x existed, exists, or will exist at all and only the times in S . The concept of exact temporal occupation makes sense both on growing block and presentism. Our method of counting future K s will be to count how many things there are that exactly temporally occupy each S in F , and then add that up across all the values of $S \in F$.

To be more precise, suppose that for some future time t in S , there is a cardinality n (zero, finite, or infinite) such that at t the following will be true:

- (16) There presently are exactly n K s that exactly temporally occupy S .

If so, then at *every* other future time t' in S , (16) will also be true. In other words, if there is cardinality n satisfying (16) at one time t in S , that same cardinality will satisfy it at other times in S . For each of the K s that will exist at t and that will exactly temporally occupy S will also exist at t' and will also exactly temporally occupy S , and *vice versa*, so the count of K s that exactly temporally occupy S cannot change between the times in S . In this case, I will define $n_K(S)$ to be equal to n , and the above argument shows that $n_K(S)$ does not depend on the choice of the future time t in S .

The other possibility is that at no time t in S is there a cardinality n satisfying (16). For example, if K is the kind *set* and S is the set of all times, then there won't be such a cardinality as there is no cardinality of the collection of all sets.²³ In that case, I will say that $n_K(S)$ is undefined.

We now have a way of counting the number of K s that will exist. If for some set S of times containing at least one future time the quantity $n_K(S)$ is undefined, we say that there is no cardinality of K s that will exist. But if $n_K(S)$ is always defined, then we say that the cardinality of the number of K s that will exist is the sum of $n_K(S)$ as S ranges over the members of F .²⁴

It is parenthetically worth noting that the above construction has another interesting use. Lewis (2004) asked whether a presentist can make claims like "There were, are or will ever be exactly n cows", and gave a complicated nested-tense translation technique for a presentist to make such claims for finite numbers n , but noted that his technique fails for infinite cardinalities n . My above technique, however, allows a presentist to do cross-time counting with both finite and infinite cardinalities. All that's needed is to drop the requirement that the sets S in F contain at least one future time, and then instead of saying that at some future time t in S the cardinality n will satisfy (16), we would say that some time t in S the cardinality n did, does, or will satisfy (16).

²³ By definition, only sets have cardinalities. But the collection of all sets is not a set, as follows from the Russell paradox. For if it were a set A , then by the Axiom of Separation we could form the Russell subset $R = \{x \in A : x \notin x\}$ of all members of A that aren't members of themselves. Then, R would be the set of all sets that aren't their own members, and so we would have $R \in R$ if and only if $R \notin R$, a contradiction.

²⁴ If we want to be more precise about it, we can let the cardinality of the number of K s equal the cardinality of the set $\{\{S\} \times n_K(S) : S \in F\}$, where we think of the count $n_K(S)$ as an ordinal.

2

Infinite Regresses

1. How to Violate Causal Finitism

Causal finitism says that there can't be infinitely many causes behind an effect (the wording is deliberately vague—we will strive for greater precision in Chapter 7). There are intuitively two ways for this to be violated:

- (a) an infinite causal regress, or
- (b) an infinite number of causes cooperating together to produce a single effect.

(Fig. 2.1.) In fact, it will turn out to be a theorem that these are the only two ways for causal finitism to be violated. Causal finitism then says that type (a) violations are impossible as are type (b) violations. In this chapter, I will focus on arguments against causal regresses, i.e., type (a) violations. I will proceed by dividing typical causal regresses into three main types, and argue against each subtype. The arguments differ in strength, but it will be simpler to assume that no causal regresses are possible than just that no causal regresses of one particular type are possible. All in all, we will have a good case against regresses and thereby some evidence for causal finitism in general. The arguments of this chapter are not focused on paradoxes, and thus adduce a different kind of evidence for causal finitism than will be given in the remainder of the book.

Given a plausible graph-theoretic formulation of causal finitism, and given a version of the Axiom of Choice from set theory, in the Appendix of this chapter it is proved that (a) and (b) are the only two ways of violating causal finitism.

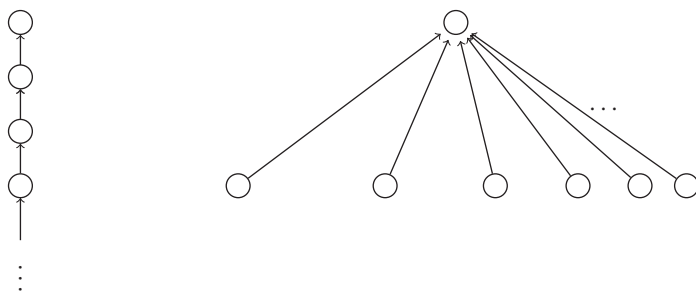


Fig. 2.1 The two ways of violating causal finitism: regress (left) and infinite cooperation (right).

Thus, violations of causal infinitism come in two varieties: infinite regresses like (b) and cases where infinitely many causes work together for a single effect. Perhaps surprisingly Thomson's Lamp, with which the book began, is of the second sort. Each flipping of the switch causally contributes to the final state of the lamp. But there is no backwards regress, because although each switch flipping is followed by an infinite number of further flippings, it is preceded by at most a finite number.

2. Infinite Causal Regresses

Infinite causal regresses appear to divide into three main types:

- (i) regresses without anything outside the regress causing the items in the regress,
- (ii) regresses with something outside the regress acting through the items in the regress as intermediate causes, and
- (iii) regresses with something outside the regress directly causing all the items in the regress.

An example of a type (i) regress is given by Hume's (1779) non-theistic infinite causal regress explanation of the universe: the present state is caused by an earlier, which is caused by a yet earlier one, and so on. Type (i) regresses are distinguished from types (ii) and (iii) by being uncaused.

A paradigmatic example of a type (ii) regress is the case of the dropping of an apple causing the landing of an apple on the ground in a universe with continuous time where an infinite number of intermediate apple-falling states act as intermediate causes. (For further discussion of such cases, see Chapter 8.)

A type (iii) regress arises in some theistic scenarios. Thomas Aquinas famously thought that although in fact God created a world with a finite past, he could have instead created a world with an infinitely long past.¹ Moreover, he thought that God was directly the primary cause of every event in the world.

Formally, other options are possible. For instance, one could imagine an infinite regress where the even-numbered items are directly caused by an outside item and the odd-numbered ones function as intermediate causes. Or one could have a regress similar to type (iii) but where only some infinite initial collection of causes is caused by the outside item. But it is plausible that if the three main types are impossible, the more recondite cases will be impossible as well. For instance, if the even-numbered items are intermediate causes, then by dropping the odd-numbered items, we get a type (ii) regress.

¹ Aquinas (1920, I.46.2) argues that one cannot prove that the past is finite, and that one shouldn't try lest one "give occasion to unbelievers to laugh". If he is right about that, this book might at least provide some comedy to the reader.

Moreover, my arguments here are only a part of my general arguments for causal finitism. The more cases of infinite causal history we can rule out, the more likely it is—think of it as an inductive argument—that *all* infinite causal histories are impossible.

We now consider the three types of regresses in order.

3. Type (i): Uncaused Regresses

3.1 *Viciousness*

Some regresses are vicious. To explain why the earth doesn't fall by saying that it stands on the back of a turtle is to leave unanswered the question of why the turtle doesn't fall. To add to the story an infinite sequence of turtles, one under the other, is to affirm a vicious regress. Likewise, *pace* infinitists like Klein (1998), to believe one proposition on the evidence of another proposition, and to believe the second because of a third, and so on *ad infinitum*, is to have a vicious justification regress.

But not all regresses are vicious. For all we know, there is today, and there will be tomorrow, and for each day there will be a next. There is the number 0, the number 1, the number 2 and so on (as well as the number -1 , the number -2 , and so on). If one has evidence for p , one has evidence that p is true, and that it's true that p is true, and that it's true that it's true that p is true, again *ad infinitum*.

There seems to be a very obvious way to divide up the cases I picked out above. The vicious cases are *dependence regresses*, regresses where each item depends on the next: the earth's position depends on the turtle and so on, the justification of one proposition depends on another, and so on. The virtuous ones are not dependence regresses; there may be dependence, but in a reversed direction (the evidence for p typically doesn't depend on the evidence for the truth of p , but *vice versa*).

Causation is a kind of dependence. Hence infinite causal regresses are vicious. And, plausibly, vicious regresses are impossible.

But we are moving too quickly. For, bracketing causal finitist considerations, it seems possible to have a non-vicious justification regress, and justification involves a kind of dependence. Let p_n be the proposition that there are at least n unicorns in the universe. For simplicity, suppose p_n is true for all n so there are infinitely many unicorns. Suppose further that there are infinitely many expert unicorn observers, and that the n th of them testifies to me that p_n is true. Then my belief that there is at least one unicorn is justified by my belief that there are at least two, which in turn is justified by my belief that there are at least three unicorns, and so on (Fig. 2.2).

This case differs from a vicious regress of justification where the justification comes in, as it were, from thin air (or from infinity?). For in our case, all the justification bottoms out in an unparadoxical way in the testimony of the observers. Still, it *is* an infinite sequence of justification relations, with p_n being justified by

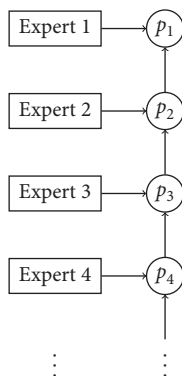


Fig. 2.2 The testimony of unicorn experts.

p_{n+1} for all n . This is a case of overdetermination at the level of the steps in the regress: p_n is justified by the n th expert's testimony as well as by p_{n+1} .

Similar overdetermination cases can be manufactured in the case of causation. Suppose that there are infinitely many lights, numbered $1, 2, \dots$, each equipped with its own switch and light detector, and set up so that the n th lamp turns on if its switch is depressed or the light detector is activated or both. Further, the light detector of the n th lamp points to the $(n + 1)$ st lamp and hence turns on the n th lamp if the $(n + 1)$ st is on. Suppose that simultaneously all the switches are pressed. Then the first lamp is on because it detects the second being on and the second is on because it detects the third being on, *ad infinitum*. But they are all also overdetermined to be on because of the switches, and there appears to be nothing vicious about the case.

Perhaps, though, in cases of overdetermination we shouldn't say that we have dependence on each of the overdetermining items. Rather, we have dependence on the disjunction of overdeterminers. If so, then we do not have regresses. For instance, the justification of p_1 depends on the justification of p_2 disjoined with the testimony of expert 2, but this disjunction does not get its justification from p_3 disjoined with the testimony of expert 3. Similarly, lamp n being on depends on the disjunctive event of lamp $n + 1$ being on or the switch of lamp n being depressed. And perhaps this disjunctive event should not be taken to depend on the disjunctive event of lamp $n + 2$ being on or the switch of lamp $n + 2$ being depressed. If this suggestion succeeds, then the case isn't in fact a case of a dependency regress.

On the other hand, if these overdetermination cases are genuinely non-vicious dependency regresses, then we should weaken the principle that all dependency regresses are vicious. Rather, we should say that those dependency regresses which do not bottom out in something outside the regress are vicious. In the causal case, type (i) regresses—uncaused regresses—will be vicious. And it is plausible that there are no vicious regresses.

3.2 *Vicious regresses and the Hume–Edwards Principle*

There is fairly robust intuition that showing that a theory posits a vicious regress is a decisive objection to the theory. That intuition is at the heart of many well-known philosophical arguments, like the Third Human Being² or the Bradley Regress (Bradley 1893). But if vicious regresses were possible, then a theory's positing such a regress shouldn't be a decisive objection to the theory. So we have reason to think that vicious regresses are impossible.

We could simply take the impossibility of vicious regresses as rock bottom, but we could also try to derive that impossibility from a Principle of Sufficient Reason (see Pruss 2006, Della Rocca 2010, and Pruss 2017 for defenses of the Principle), on the grounds that vicious regresses involve something being unexplained, namely why the whole regress obtains—why, say, the whole infinite pillar of turtles supporting the earth doesn't fall. This fits particularly well with the intuition that vicious regresses are ones that do not bottom out in something³ outside the regress. For it is precisely when given a regress that does not bottom out in this way, there is no explanation of the regress as a whole.

There is, however, an objection to this line of thought. Famously, Hume (1779) thought that a causal regress would be self-explanatory. He claimed that if every item in a collection was causally explained, the whole collection would be explained. In the literature, this is called the Hume–Edwards⁴ Principle (HEP), and at first sight it seems very plausible. But then in a regress, every item is explained by a preceding item, so Hume concluded that the whole collection is thereby explained.

One can argue, however, that HEP is false in the case of infinite causal regresses. To see this, suppose time is dense—between any two times, there is another time. Now imagine that a particle comes into existence from some cause in such a way that at every time strictly after noon the particle exists, but at noon and at earlier times the particle does not exist. The existence of the particle at any time t later than noon can be causally explained by means of the existence of the particle at an intermediate time t' between noon and t and the particle's propensity to maintain its existence. But it would be absurd to conclude that this particle's existence at all times after noon has been explained by a story that never mentioned the cause at noon (cf. Pruss 1998).

We now have two options with regard to HEP. The first, and I think more plausible, option is to say that the initial plausibility of HEP relied on one not having thought enough about infinite cases. In a finite case, if each item in a collection has a causal explanation, then there will be at least one explanation outside the collection, on pain of explanatory circularity. It is then plausible that all such external explanations, taken

² Plato (1996, 132a–b) gives a version involving a regress of largenesses, rather than humanities.

³ Or a plurality of things. In our case of the lamp regress of light detectors with overdetermining switches, the regress does bottom out in the plurality of switches.

⁴ See Edwards (1959).

together, will explain the whole collection. However, in an infinite case, there could be purely internal explanations.⁵ And it is just not plausible that there we would have an explanation of the whole collection.

There is also a bolder move one could make instead. HEP is initially plausible. We *can* retain it without any qualification as long as we deny the possibility of infinite regresses. Thus we have an argument, from HEP plus the fact that infinite regresses are not explanatory, for the impossibility of infinite regresses. But this argument is not very compelling, as HEP is not very compelling in infinite cases.

3.3 Regresses and explanatory loops

Explanatory loops are intuitively problematic, and a number of authors (e.g., Dummett 1986, Pruss 1998, and Meyer 2012) have noted similarities between explanatory loops and infinite regresses (though in the case of Meyer 2012, in order to defend explanatory loops).

Suppose there have always been chickens and eggs, and each chicken came from an egg while each egg came from a chicken. Then the plurality of the eggs is explanatorily prior to the plurality of the chickens, say, because of this plausible principle:

- (1) If each of the *ys* has at least one of the *xs* explanatorily prior to it, then the *xs* are explanatorily prior to the *ys*.

But for exactly the same reason, the plurality of the chickens is explanatorily prior to the plurality of the eggs.

Thus we have circularity in explanatory priority: the chickens are explanatorily prior to the eggs and the eggs are explanatorily prior to the chickens. Now, plausibly:

- (2) It is not possible to have a circularity in the order of explanation when the type of explanation is kept fixed.⁶

The chickens and eggs case violates (2), since in both directions the explanation is provided by efficient causation. Hence the backwards infinite regress of chickens and eggs is impossible.

But the same is true for *any* backwards infinite regress. If a_1 is caused by a_2 which is caused by a_3 and so on, then we can label the odd-numbered entries “chickens”

⁵ The existence of a crucial difference between finite and infinite cases seems to have been first noticed by Rowe (1970).

⁶ I doubt the proviso is needed, but some think it is. For instance, on Humean views of laws, particular events ground the laws, but the laws nomically explain the particular events. This makes for an explanatory circularity, but the types of explanation in the two directions are different—grounding and nomic subsumption, or metaphysical and scientific (Hicks and van Elswyk 2015). Those who think the proviso is not needed, will think this is a refutation of Humeanism. Humeans will think this is an argument in favor of the proviso.

and the even-numbered ones “eggs”, and the argument goes through.⁷ Hence, infinite causal regresses are impossible.

The most problematic part of the argument is (1). For instance, suppose the *ys* are all of Charlemagne’s descendants, and the *xs* are Charlemagne and all his descendants. Then each of the *ys* has an *x* explanatorily prior to it, namely Charlemagne. But while it is clearly correct to say that Charlemagne is explanatorily prior to his descendants, it is not clearly correct to claim that the larger plurality is prior to the smaller.

One can rule out the Charlemagne plus descendants counterexample by adding to (1) the condition that the *xs* and *ys* have no entities in common. That will still be enough to carry our application to the chickens and eggs regress. But now consider this case: *A* causes *B* which causes *C*. Let the *xs* be *A* and *C*, and let the *ys* consist of just *B*. Once again, it is not clearly correct to say that the *xs* are prior to the *ys*: one of the *xs* is not prior to any of the *ys*. The following refinement, however, takes care of this case as well:

- (3) Suppose (a) each of the *ys* has at least one of the *xs* explanatorily prior to it, (b) the *xs* and *ys* have no entities in common, and (c) each of the *xs* is prior to at least one of the *ys*. Then the *xs* are explanatorily prior to the *ys*.

However, the replacement of (1) with (3) fails to provide a direct argument against the backwards infinite chickens and eggs regress. For suppose that Little is the last of the chickens, and has no eggs. Then (3) does not allow us to conclude that the chickens are prior to the eggs, since one of the chickens, namely Little, is prior to none of the eggs.

Nonetheless, the argument can be rescued. For, surely:

- (4) If a backwards infinite regress of chickens and eggs is possible, then a bidirectional infinite regress of chickens and eggs is possible as well.

It is, after all, surely possible that in addition to the backwards regress, Little has an egg, and that egg hatches into a chicken, and so on. But if we have such a bidirectional infinite regress of chickens and eggs, then every egg has a chicken prior to it and every chicken is prior to some egg, so the chickens are prior to the eggs, and similarly the eggs are prior to the chickens, in violation of (2). Thus, (3) and (4) together yield an argument against backwards infinite regresses.

While there is some force to this anti-regress argument, that force is diminished by the need to have two provisos in (3) in addition to the basic condition (3)(a), as (3) sounds somewhat *ad hoc*.

This argument appears to be against *all* types of regresses. But it is most telling against uncaused regresses, like the uncreated backwards chain of chickens and eggs. For what seems most objectionable about causal loops is the attempt to lift oneself

⁷ This argument is based on ideas in Pruss (1998).

by one's bootstraps. A causal loop with an overarching cause outside the loop seems less problematic.

4. Type (ii): Causation Passing through Infinitely Many Steps

One kind of infinite regress of causes is an infinite regress of intermediate causes between an initial cause and a final effect. The apple's landing on the ground is caused by the apple's breaking off from the branch. But on a plausible interpretation of Newtonian physics—though we will consider alternatives in Chapter 8—in between there are infinitely many intermediate causes: the apple's being half-way down, the apple's being a quarter of the way down, and so on. In this section I will offer a metaphysical argument against such cases coming from ideas by Robert Koons.⁸

Some instances of causation are derivative. If I press a button which turns on the lights and alerts the burglars, pressing the button causes the burglars to be alerted. But this is a derivative instance of causation: the button press causes the alerting of the burglars *by* causing the light to go on and *by* the light causing the burglars to be alerted.

We can then offer this plausible derivation principle:

- (5) Derivative instances of causation must ultimately be grounded in fundamental instances of causation.

Principle (5) creates a serious problem for stories like the one about the apple. For the apple's being half-way down is an intermediate cause of its being on the ground, and so the apple's breaking off from the branch doesn't fundamentally cause the apple being on the ground. Indeed, if between any two items in the sequence of apple states there is another, there won't be *any* fundamental cases of causation in the sequence. Each case of causation in the sequence will go through an intermediate cause. And it is implausible to think that the causation within the sequence will be grounded in some fundamental causation outside the sequence. Thus, the story of the apple violates principle (5), as do other dense causal sequences—sequences where between each pair of items there is an intermediate cause.

More generally, scenarios where something causes an effect through a backwards-infinite causal sequence of intermediate causes violate the derivation principle. For suppose that we have *a* causing *e* through a backwards infinite sequence $\dots, c_{-3}, c_{-2}, c_{-1}$ of causes leading to $c_0 = e$. Then *a*'s causing c_n is derivative from *a*'s causing c_{n-1} and c_{n-1} 's causing c_n , and so there is nothing in the sequence that *a* causes non-derivatively.

⁸ I heard these ideas in our joint seminar on neo-Aristotelian metaphysics in the fall of 2008.

The only way to reconcile this kind of a sequence with the derivation principle and the intuition that causation through an intermediate cause is derivative appears to be to suppose that a causes *non-derivatively*, and hence directly, all of the c_{-n} . Thus, c_{-n} will be directly caused *both* by a and by $c_{-(n+1)}$. Moreover, $c_{-(n+1)}$ must be extraneous to the causation of c_{-n} by a , since otherwise a 's causing of c_{-n} will be derivative. Thus, this is a case of overdetermination of c_{-n} by both a and $c_{-(n+1)}$, and is a type (iii) regress.

Besides the intuitive plausibility of the derivation principle, a reason to accept it is that if we deny the derivation principle then we get an infinite regress of metaphysical grounding: a derivative case of causation grounded in one or more other derivative cases of causation, *ad infinitum*. Moreover, if we think that causation is not grounded in phenomena other than causation, there won't be anything outside of *this* grounding regress that grounds all the items in the regress. Thus, interestingly, the derivation principle lets us show that a causal regress *with* a first element implies a truly vicious grounding regress *without* a first element.

5. Type (iii): Outside Cause Directly Causing Each Item

5.1 Options

This leaves us with one last main type of regress. This is a case where there is an outside cause that directly causes every item in the regress. Two examples of this kind of a regress have been mentioned so far. The first was a non-causal regress involving justificatory overdetermination. The second was a theistic suggestion that God could create a universe with an infinite past, with God directly causing each item in the world.

Let us consider the theistic model (Fig. 2.3). There will be an infinite sequence of created causes $\dots, c_{-3}, c_{-2}, c_{-1}$, with God being directly the cause of each c_{-n} . So, c_{-n}

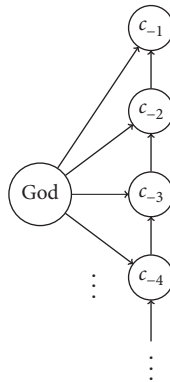


Fig. 2.3 A theistic non-vicious regress?

has two apparently immediate causes, God and $c_{-(n+1)}$. For simplicity, I will take $c_{-(n+1)}$ to include in itself all the creatures that are working together to produce c_{-n} (thus, the items in the regress might be pluralities or mereological sums).

There are now four options, depending on which, if any, of the two causes is causally sufficient to explain c_{-n} . Being sufficient to causally explain is not the same as being causally sufficient. For instance, a wavefunction and an act of measurement may be sufficient to causally explain why an electron was detected in a particular location x , but in indeterministic Quantum Mechanics the wavefunction and act of measurement are not causally sufficient for that particular effect, since that very wavefunction and act of measurement could have resulted in the electron being detected in a location other than x (causal sufficiency, as I use the phrase, implies that the cause determines the effect). On the other hand, the wavefunction on its own (or the act of measurement on its own) is not even sufficient to causally explain the outcome.

Suppose God's causal activity is sufficient to causally explain c_{-n} and $c_{-(n+1)}$ as well. Then this is basically a case of overdetermination. On this picture, we have something structurally similar to the justificatory overdetermination regress we discussed in Section 3.1. Here, we need to understand "overdetermination" in an extended sense, compatibly with the causation being indeterministic. Thus, if one random process has a chance of resulting in Smith's fatal poisoning and another has a chance of resulting in Smith's strangulation, and both processes go off and cause their effects, it might be that Smith's death was "overdetermined" by the two processes, because each process is sufficient to causally explain Smith's death, even though neither process determines the outcome as neither process is a sufficient cause (nor are the two together sufficient).

Next, suppose that God (or His causal activity) is sufficient to causally explain c_{-n} but $c_{-(n+1)}$ is not. If $c_{-(n+1)}$ is to play any role here, it must be as *part* of an overdetermining (in the above weak sense) cause of c_{-n} , since God is sufficient to causally explain c_{-n} . Thus there will still be overdetermination: there will be some d (perhaps a plurality) such that God is sufficient to causally explain c_{-n} , and $c_{-(n+1)}$ plus d are sufficient to causally explain it as well. (Note that d could even be God.)

Similarly, if $c_{-(n+1)}$ is sufficient to causally explain c_{-n} but God is not, there will also be some sort of overdetermination here.

That leaves one last option, where neither $c_{-(n+1)}$ nor God is sufficient to causally explain c_{-n} , but they work together to explain c_{-n} . Here we have a vicious regress: c_{-1} is caused by God and c_{-2} , and c_{-2} is caused by God and c_{-3} , and so on. In this regress only one of the two causes in the regress has a cause in the previous level, since God is uncaused. However, it is nonetheless a vicious regress. The case is akin to this one: Imagine an infinite floor, and suppose Jim spilled an infinite container of oil on it. On the floor there are infinitely many people in motion, with the first one moving because the second impacted her on a floor Jim spilled oil on, the second impacting the first because the third impacted her on a floor Jim spilled oil on, the third impacting the second because the fourth impacted her on a floor Jim spilled oil

on, etc. This is just as vicious a regress of impacts as it would be if the people were all floating in an oil-less vacuum, even though only one of the two items at each level of the regress (namely, the movement of a person) is caused by the previous level of the regress. The reasons to deny uncaused regresses apply just as much to these kinds of regresses.

5.2 *Regresses with outside overdetermination*

Hence, type (iii) regresses divide into two options: we either have (iii-a) a case of something like overdetermination (albeit perhaps of an indeterministic sort) or (iii-b) a vicious regress of a non-overdetermining structure. Neither kind of regress is much discussed in the literature.

At this point one must note a point of weakness in the argument of this chapter against causal regresses. I do not have a very compelling direct argument against the overdetermination cases of infinite regresses. These are regresses where each item has an ultimate causal explanation, namely in terms of the outside cause, which along with the items in the regress overdetermines the effect. There is a sense in which these regresses are not vicious. Nonetheless, there is some intuitive reason to think that if overdetermining infinite causal regresses are possible, so are non-overdetermining ones. And the overall argument of this book for causal finitism as a fairly simple explanation of what goes wrong in all the causal paradoxes of infinity gives one reason to deny the possibility of these cases as well.

However, even if there is no *very* compelling argument against the overdetermination scenario, there is an intuitive reason to be skeptical of it. For suppose that we have a causal regress $\dots, c_{-3}, c_{-2}, c_{-1}$, with each item causing the next, and suppose that there is an item d outside the regress such that d directly causes each of the c_{-n} , and c_{-n} is overdetermined by d and $c_{-(n+1)}$ (the picture will look like Fig. 2.3, but with d in place of God). If this is possible, it should also be possible to have a scenario like the above where the above-mentioned causal relations between the c_{-n} and between d and the c_{-n} are all the causal relations there are. Let's suppose that.

Then each instance of d 's causal influence is overdetermined by d and by something other than d . But something that causes only in an overdetermined way is causally otiose, and it should be possible to remove causally otiose influences. Thus, intuitively, there should be a possible world just like the one described but without d , or (if d is a necessary being) without d having any relevant causal role. In such a world, we have an uncaused causal regress. But that is a type (i) regress, and there is reason to think such are impossible.

There is one technicality here. Given essentiality of origins, there is some reason to think that if one removed d 's causal influence, the c_{-n} could no longer exist. If so, then we should say that there is a possible world with a causal regress of c_{-n}^* which are much like the c_{-n} except that their causal origins don't include d . (Compare: Arguably, Socrates couldn't exist if his parents didn't, but someone much like Socrates could exist instead.) And that is all that is needed for the argument.

This is not a very strong argument. By removing d 's influence from the world one removes that which makes the regress non-vicious, and that gives one reason to doubt the possibility of removing d . Still, the intuition about the removability of something all of whose influence is overdetermined has some force, and if one combines this argument with the indirect argument that it is more simple and elegant to uniformly deny all infinite causal regresses we get a strong consideration.

6. *Analogy with Axiom of Regularity

The Axiom of Regularity in set theory is a standard part of the Zermelo–Fraenkel (ZF) axioms. Formally, the axiom says that every set a contains a member b such that no member of b is a member of a (i.e., $a \cap b = \emptyset$). It follows, for instance, that no set can be a member of itself (for if $a \in a$, then the unique member of $\{a\}$ has an element in common with $\{a\}$, namely a), as well as that there is no set of all sets (since it would be a member of itself). Furthermore, Regularity rules out infinite membership regresses $\dots \in a_{-3} \in a_{-2} \in a_{-1} \in a_0$. For if we had such a regress, then we could form the set $a = \{a_n : n \leq 0\}$, and observe that every member a_n of a has a member in common with a , namely a_{n-1} , contrary to Regularity.

The intuition behind Regularity appears to be based on the idea that a set is grounded (at least partially) in its members. Suppose now that a set a is a counterexample to Regularity. Then each member b of a has one of two properties: either b has itself as a member or b has some other member of a as a member. It is absurd to suppose that something could be self-grounded—that would be like being self-caused.⁹ That leaves the other option, that every member of a is grounded in at least one other member of a . The intuition behind Regularity holds that this kind of grounding relationship is absurd, at least in the case where the grounding relationship is constituted by set membership.

But arguably there is nothing special about set membership here. It seems plausible that just as we should rule out a set of objects each of which is set-membership-grounded in at least one of the others, we should rule out all cases of a set of objects each of which is grounded (even only partially) in one of the others. But to do that is to rule out a grounding regress of objects, $\dots, x_{-3}, x_{-2}, x_{-1}, x_0$, where x_{n-1} grounds x_n for all $n \leq 0$. For if we have such a grounding regress, then the set $\{x_n : n \leq 0\}$ has the property that every member of that set is grounded in some other member since x_n is grounded in x_{n-1} .

It is further plausible to think that there is at least an analogy between causation and grounding (Schaffer 2016). Given this, reasons to reject grounding regresses by analogy yield reasons (perhaps weaker ones) to reject causal regresses.

⁹ Descartes thought that God was self-caused, though this was not a mainstream view in the Western monotheistic tradition. It is a mainstream view that God is *a se*, but it is reasonable to take this to express ontological independence rather than self-grounding.

7. Evaluation

A violation of causal finitism must include an infinite causal regress or an infinite number of causal cooperators or both. A number of the paradoxes we will consider in subsequent chapters will involve infinite causal cooperation, while the focus of this chapter was on infinite causal regresses.

Paradigmatic examples of causal regresses come in three types. First, we have the uncaused infinite regress, akin to Hume's atheistic explanation of the universe by an infinite regress of causes. It is plausible to reject *vicious* regresses as impossible, and uncaused causal regresses are vicious.

Second, we have an infinite regress—in fact, paradigmatically, a dense causal sequence—of instrumental causes, with a first cause that acts through these instrumental causes. Such a regress is incompatible with the idea that causal relations are grounded in fundamental causal relations.

Third, we have the most complicated case. Here we have an infinite regress and an item outside the regress that causes every item of the regress. Some versions of this situation involve an explanatory viciousness as in the first type of regress, and can be argued against in a similar way. But there is a particularly pesky version involving overdetermination, where each item in the regress is overdetermined by an outside cause *and* the preceding item in the regress. Here the argument is weakest: I suggest that there is good reason to have a uniform view of all the regresses and to hold that they are all impossible. This argument will become stronger in subsequent chapters, however, as we come up with more and more problematic situations involving infinite causal histories—regressive or not—thereby giving us good reason to reject the possibility of *all* infinite causal histories.

Appendix: *Two Kinds of Violations of Causal Finitism

Say that a *causal nexus* is a directed graph whose nodes are causal relata, with the nodes being joined with arrows corresponding to relations of partial causation or causal contribution (we leave open the possibility of a node being connected with infinitely many nodes). If c and d are two nodes in a causal nexus, say that $c < d$ if and only if there is a finite sequence of arrows from c to d all pointing in the same direction, i.e., $c = c_1 \rightarrow c_2 \rightarrow \dots \rightarrow c_n = d$.

We now say that a *history* of an item e is a causal nexus H that contains e as a node and is such that every other node c of H satisfies $c < e$. A causal nexus is *realized* provided that all the nodes of the nexus really exist and the arrows correspond to really obtaining relations of partial causation. A causal nexus is *possible* provided that it is metaphysically possible for it to be realized.

Causal finitism can be taken to be a claim that every possible history has only finitely many nodes.¹⁰

¹⁰ In Chapter 7, Section 5.1 we will consider a generalization of this.

The Axiom of Dependent Choice is a particularly plausible weaker version of the Axiom of Choice. It holds that if we have a relation R on some set S such that for each $x \in S$ there is a $y \in S$ with xRy , then there is a countably infinite sequence x_1, x_2, \dots such that $x_n R x_{n+1}$ for all n . Intuitively, you choose the first element of the sequence, then the next, and so on.

THEOREM. *Assume the Axiom of Dependent Choice. Then a history H of an item has infinitely many nodes if and only if:*

- (a) *it has a node that has arrows pointing to it from infinitely many nodes, or*
- (b) *it has an infinite backwards sequence $\dots \longrightarrow c_3 \longrightarrow c_2 \longrightarrow c_1 \longrightarrow c_0$ where all the c_n are distinct.*

The conditions (a) and (b) are not, of course, exclusive. And it is clear that if either condition holds, then H has infinitely many nodes. To prove the converse, suppose that H has infinitely many nodes but (a) is false. Thus every node of H has at most finitely many arrows pointing to it. To complete the proof we must show that there is a regress as in (b).

Say that a *path* from node a to node b in a directed graph is a finite sequence of nodes x_1, \dots, x_n such that $a = x_1$ and $b = x_n$ and there is an arrow from x_i to x_{i+1} for all i (the trivial case is where $x_1 = x_n = a = b$). Then the *length* of the path is $n - 1$. Suppose that the history H is a history of e . Then let the *depth* $d(a)$ of a node a be the length of the shortest path from a to e (Fig. 2.4).

Say that an *ancestor* of a node b is any node $a \neq b$ from which there is a path to b of strictly decreasing depth, i.e., a path $a = x_1 \longrightarrow \dots \longrightarrow x_n = b$ such that $d(x_i) > d(x_{i+1})$ for all i (hence $d(x_i) = d(b) + n - i$, since the depth of successive nodes in a path can only differ by one). If further $d(a) = d(b) + 1$, then I will say that a is a *parent* of b .

Next let $I_n = \{a : d(a) = n\}$ for $n \geq 0$. Observe that each set I_n is finite (this can be proved by induction using the fact that a finite union of finite sets is finite and each node has only finitely many arrows pointing to it by the denial of (a)), that the I_n are disjoint, and that their union is all of the nodes of H . Let $L_n = \{a : d(a) \leq n\}$ and $M_n = \{a : d(a) \geq n\}$. Then $L_n \cup M_{n+1}$ is all the nodes of H for any n , and L_n is finite, being a union of the $n + 1$ finite sets I_i for $0 \leq i \leq n$.

Finally, let N be the set of those nodes of H that have infinitely many ancestors. I first claim that $e \in N$, so that N is non-empty. For every node c of H other than e has a path to e as H is a history. A path from c to e of minimal length will then exist, and it is easy to see that that path will have strictly decreasing depth, so that c is an ancestor of e . But there are infinitely many nodes c other than e in H (since we assumed H has infinitely many nodes), so $e \in N$.

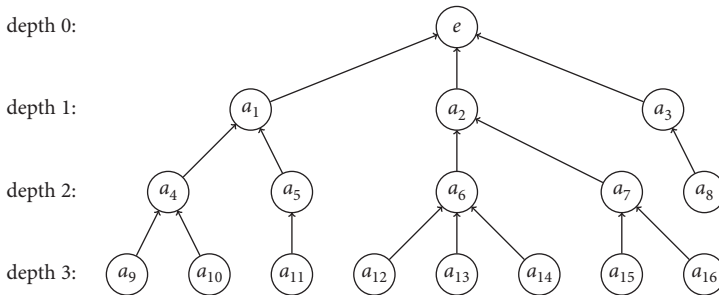


Fig. 2.4 Here, $I_1 = \{a_1, a_2, a_3\}$, $L_1 = \{e, a_1, a_2, a_3\}$, and $M_2 = \{a_4, a_5, \dots\}$.

I now claim that for any any $n \geq 0$ and any $b \in I_n \cap N$, there is a node $a \in I_{n+1} \cap N$ such that H has an arrow from a to b . For b has infinitely many ancestors, and hence infinitely many ancestors in M_{n+2} (since L_{n+1} is finite and $L_{n+1} \cup M_{n+2}$ is all the nodes of H). Each ancestor has a path to b of strictly decreasing depth. That path must cross the finite set I_{n+1} immediately before ending at b . Thus b has infinitely many ancestors each of which is an ancestor of a parent of b . Since b has only finitely many parents, as each node has only finitely many arrows going to it, it follows that at least one of b 's parents has infinitely many ancestors, and hence at least one of b 's parents is in N . But all of b 's parents are in I_{n+1} , so the proof of the claim is complete.

The claim shows that every member of N has at least one parent in N . Now define the relation R on N by xRy if and only if y is a parent of x . This relation satisfies the conditions for the Axiom of Dependent Choice, so there is an infinite sequence c_1, c_2, \dots such that $c_n R c_{n+1}$, i.e., c_{n+1} is a parent of c_n for all n . This is the desired infinite regress (since the depths of all the members of the sequence are different, the elements in the sequence are different).¹¹

¹¹ This proof is based on suggestions by Will Brian and Daniel Herden on how to prove König's Lemma using the Axiom of Dependent Choice.

3

Supertasks and Deterministic Paradoxes

1. Introduction

In this chapter we consider several paradoxes centered around deterministic processes (though occasionally considering an indeterministic variant). The first, Thomson's Lamp, was already introduced in Chapter 1, but now will be discussed in more detail. It will give some evidence for causal finitism, but not much. Next, the Grim Reaper paradox will be considered, which will be a lot more compelling. Both paradoxes involve supertasks: tasks that have an infinite number of events happening within a finite interval. Then a Newtonian argument *against* causal finitism will be given, but we will see that Newtonian considerations on the whole *support* causal finitism (Newtonian physics is false, of course, but the arguments will only use the metaphysical possibility of such physics). Finally, there will be a brief discussion of the paradox of an eternal life after an eternal life, and parallels will be adduced between the paradoxes of this chapter and the Grandfather paradox argument against causal loops and time travel.

2. Thomson's Lamp Revisited

2.1 Introduction

Recall that Thomson's Lamp has a toggle switch which swaps its state between on and off. At 10:00 am the lamp is off and its switch is toggled countably infinitely often between 10:00 am and 11:00 am. In the usual version of the scenario, the toggglings bunch up at 11:00, occurring, say, at 10:30, 10:45, 10:52.5, and so on. Nothing else changes the state of the lamp. This is a *supertask*: a scenario where infinitely many things are done in a finite amount of time.

Paradox now ensues when we ask whether the lamp is on or off at 11:00. Neither answer seems satisfactory.

2.2 Causal finitism

As noted in Chapter 1, causal finitism gives a neat resolution to the paradox. Since infinitely many causes cannot bear on one thing—the lamp's state at 11:00—the

scenario is simply impossible. To transpose this point to an argument for causal finitism, say that:

- (1) If it's possible for infinitely many causes to bear on one thing, it is possible for them to bear on that thing in the way that the Thomson's Lamp story says they do.

But:

- (2) It is not possible for causes to be arranged as in the Thomson's Lamp story.
- (3) So, it is not possible for infinitely many causes to bear on one thing; i.e., causal finitism is true.

However, we need to consider alternative responses to the paradox.

2.3 Non-standard analysis

One would like to say: Whether the lamp is on or off at 11:00 depends on whether "infinity is even or odd".¹ But it makes no sense to say that infinity is even or odd. After all, $\infty = \infty + \infty$, which would suggest that ∞ is even, but also $\infty = \infty + \infty + 1$, which would suggest that ∞ is odd.

But this was too quick. We can do better with infinities. Non-standard analysis is a mathematically rigorous way of handling arithmetic with infinite and infinitesimal quantities by extending the real numbers to the "hyperreal" numbers, which are a set-theoretic construction out of the reals.² There is even a transfer principle guaranteeing hyperreal quantities obey analogues of all the standard rules of arithmetic, provided that one takes care to translate concepts accordingly. In particular, if N is infinite, it is *false* that $N = N + N$, since by the ordinary rules of arithmetic, if N is non-zero, we can divide both sides by N and get the falsehood $1 = 1 + 1$.

The ordinary real numbers \mathbb{R} are Archimedean in the sense that for every real number r there exists an integer n such that $r < n$. But the hyperreals ${}^*\mathbb{R}$ are not Archimedean, since they contain infinite values. This may seem to violate the transfer principle, but it does not. For it's still true that for every hyperreal number r there exists a *hyperinteger* n such that $r < n$. The hyperintegers are a subset of the hyperreals with properties such as these: any integer is a hyperinteger; the sum, difference, and product of any two hyperintegers is a hyperinteger, and hyperinteger arithmetic obeys all the same rules as integer arithmetic; and for every hyperreal r , there exists a unique hyperinteger m such that $m \leq r < m + 1$. If r is a positive infinity, then of course a hyperinteger n such that $r < n$ will be an *infinite* hyperinteger. So the claim that the real numbers are Archimedean does have an analogue in the hyperreal context, but that analogue only says that the hyperreals are "hyper-Archimedean".

¹ My son said something like this when I first told him the paradox.

² *The construction depends on the Axiom of Choice, or at least on a special case of the Boolean Prime Ideal axiom.

Given the hyperreals, we can try to resolve Thomson's Lamp by saying that the number of switch togglings will be some infinite hyperinteger N , and the lamp will be on or off depending on whether N is even or odd.

Indeed, it now does make sense to say whether N is even or odd. But there is another problem. For on the usual setup the n th toggling occurs at $60 - 60/2^n$ minutes after ten o'clock. If N is an infinite hyperinteger such that there are N togglings, then the N th toggling would occur $60 - 60/2^N$ minutes after ten o'clock. Since N is infinite, $60/2^N$ is infinitesimal, where a number α is infinitesimal provided that $\alpha \neq 0$ and $|\alpha| < r$ for every positive real number r . But in our paradox there were no togglings at infinitesimal moments before 11:00. Each toggling was specified to occur a certain definite, positive, *real*-numbered amount of time before 11:00.

Furthermore, the number of hyperintegers between 1 and N , where N is infinite, is guaranteed to be uncountable—these hyperintegers cannot be put in one-to-one correspondence with the natural numbers (Pruss 2014, Appendix). But we specified that there were *countably* many button togglings.

The non-standard analysis resolution, thus, fails for the normal version of Thomson's Lamp. It might provide an answer for an alternate version on which there is a hypernatural number of button togglings most of which are an infinitesimal amount of time before 11:00. But one does not resolve a paradox simply by giving a solution for a different paradox.

2.4 Special Relativity

The switch-toggling in Thomson's Lamp need to happen faster and faster without bound: the amount of time available for the switch to be toggled is the temporal spacing between togglings, and that gets smaller and smaller, converging to zero in the limit. If the switch needs to be moved a fixed amount in each case, this means that the speed of switch movement goes up without bound—and in particular eventually exceeds the speed of light, in violation of Special Relativity. We thus have good independent reason to reject the Thomson's Lamp story, regardless of any hypotheses like causal finitism.

This is correct, but it's only good reason to reject the claim that the Thomson's Lamp story is *true*. It's not a good reason to reject the claim that the story is metaphysically possible. A world governed by Newtonian physics, without any absolute speed limit, certainly seems metaphysically possible.³

Moreover, we can imagine scenarios very similar to Thomson's Lamp where there are no violations of an absolute speed limit. For instance, we might imagine a world where there are θ particles which have the property that when two of them come

³ It might be that in that world, there would be no *light*: it could be that the nature of light is tied to the relativistic laws governing it. But even if in that world there was no light, and hence no lamps, there could be something functionally behaving much like a lamp, and that's all that's needed for the paradox.

together they are invariably annihilated in a burst of energy. We could then imagine that at 10:00 am a sticky target has no particle on it, and a particle emitter moves gradually closer to the target between 10:00 am and 11:00 am, shooting off a θ particle at 10:30, 10:45, 10:52.5, and so on, at sufficient sub-light velocity that each particle can reach the target before the next one is shot without exceeding the speed of light.⁴ The first particle sticks to the target. When the second hits, both it and the first are annihilated. Then the third sticks. As a result, we have an alternation between the target containing a θ particle and the target not containing a particle, much as in the original paradox.

I shall in the future leave such fairly easy modifications of paradoxes to the reader and hence not worry much about relativistic objections.

2.5 Benacerraf's solution and the Principle of Sufficient Reason

Benacerraf (1962) argues that there just is no absurdity whether or not the lamp is on or off at the end of the experiment. Both outcomes are compatible with the story as given. Neither gives rise to a contradiction. The story doesn't determine which of the two outcomes will happen, but underdetermination is no paradox.

This is the start of a resolution, but it is not a complete resolution. After all, there is at least *some* reason to believe the Principle of Sufficient Reason (PSR)⁵ which holds that every contingent fact has an explanation. But in Benacerraf's solution nothing explains why the lamp has the state it does at the end of the scenario. A solution that requires denying the PSR has some cost.

Depending on how we read the original story, perhaps we can find a solution compatible with the PSR, though. Suppose that aliens come and instantaneously disconnect the lamp right at 11:00 in such a way that the lamp is off then (or instantaneously power up the bulb, if we prefer). In that case, there is no violation of the PSR, and yet this may be compatible with my original story.

Whether the alien case is compatible with the original story depends on how we read the statement that nothing but a button toggling changes the lamp's state. For it is not clear that the aliens' activity counts as a change of the lamp's state. For the aliens didn't come to find a lamp that was on and instead make it be off. Indeed, for every time interval at which the lamp was on, there was a later time interval at which the lamp was off, simply because of button toggling and not because of alien meddling. And then right at 11:00, the aliens came and disabled the lamp.

However, the proponent of a paradox has an easier task than the resolver. As we noted in subsection 2.3, one can't resolve a paradox by changing its story in a relevant

⁴ For instance, suppose the initial distance to the target is one meter, and the emitter moves towards the target at one meter per hour, while the particles shoot out at four meters per hour. Then at 10:30, the emitter is half a meter away from the target, and the particle will reach the target in 7.5 minutes, well before the next emission at 10:45. At that next emission, the emitter will be a quarter of a meter away from the target, and the particle will reach the target in 3.75 minutes, which is well before the next emission.

⁵ See, for instance, Pruss (2006), Della Rocca (2010), and Pruss (2017) for defenses.

way. But the proponent of a paradox is free to qualify the paradox in additional ways to make it more paradoxical. In this case, we simply stipulate that nothing besides the toggling causally affects the state of the lamp, and we rule out the aliens regardless of having to make any decision on the semantics of “change”. (And the intuitions behind the crucial premise (1) should remain.)

But perhaps there is another way of resolving the paradox without denying the PSR. It seems compatible even with the revised story that there is a law of nature that whenever a lamp has been toggled precisely at at 10:30, 10:45, 10:52.5, and so on, then its state at 11:00 is the opposite of its state before 10:30. A law of nature can do the aliens’ work. This scenario shows that the revised (and original) lamp story does not automatically violate the PSR.

Whether this scenario is possible depends on whether one can have a law of nature non-causally determining the state of the lamp at 11:00. On some Aristotelian views, for instance, laws of nature are simply statements about the regularities in the arrangement of causal powers in nature, and the explanation-giving force of laws of nature is derivative from the causal activity of these powers. The revised scenario, by limiting causal influences to button toggles, rules out a law-based explanation as well.

But now the paradox gets quite top-heavy. To respond to Benacerraf, the paradox-propounder has had to suppose one controversial metaphysical thesis, the PSR. Now she had to further rely on a controversial thesis about laws of nature. Instead of being a paradox, the story could be seen as evidence against the conjunction of the PSR with the Aristotelian view of laws. Admittedly, there is still *some* cost to Benacerraf if he has to reject at least one of the two controversial theses, but the cost is not so great.

Furthermore, the more conditions we add to the paradox, the harder it is to argue for something like (1), namely that if infinitely many causes can affect something, they can do so in the way the paradox says they do.

2.6 *Two counterfactuals*

There is, however, another way to think of the paradox, without reference to the Principle of Sufficient Reason. Start with two thoughts. If you take the situation described and shift the time of one of the button presses, while keeping the order of button presses unchanged, say moving the 10:15 button press to 10:12 or 10:17, this should not change the outcome at 11:00. The causal contribution of a button press only depends on the temporal position of the button press in the sequence of button presses, rather than the exact time:

- (4) For any shifted sequence of times of button presses within the 10:00 to 11:00 (both non-inclusive) period that keeps the mutual order the same, implementing that sequence in place of the actual one wouldn’t have affected the state of the lamp at 11:00.

The second thought is this:

- (5) For any button press in the actual sequence, removing that button press would have swapped the state of the lamp at 11:00.

But (4) and (5) cannot both be true. For imagine shifting the times of the button presses as follows: the 10:30 button press is shifted to 10:45, the 10:45 to 10:52.5, the 10:52.5 to 10:56.25, and so on. By (4), this wouldn't change the state of the lamp at 11:00. But this shift has the net effect of removing the 10:30 button press, and by (5), this would flip the outcome at 11:00. This shift can be described in two ways: in one way as a shift that satisfies the conditions in (4) and in another as a removal of one particular button press. And so the two counterfactuals give opposite results, despite having equivalent antecedents. But the only way we can have both:

- (6) If p were to hold, q would hold

and

- (7) If p' were to hold, r would hold

where q and r are logically contradictory while p and p' are logically equivalent is if p (and hence p') is impossible. But if the Thomson's Lamp story is possible, so is this shift of button press times.

Nonetheless, even this counterfactually based argument is not very convincing. The original story seems compatible, say, with there being a law of nature according to which whenever buttons are pressed at 10:30, 10:45, 10:52.5, and so on, the result is that the lamp is on at 11:00, but when they are pressed at 10:45, 10:52.5 and 10:56.25, the lamp is off at 11:00. In this case, (4) is just false: there are ways of shifting button presses without changing their order that don't change the result—and, parenthetically, the “without changing their order” isn't doing any work unless the button presses have something else to distinguish them besides their temporal position. Alternately, there could be a law of nature that guarantees that *every* countably infinite sequence of button presses between 10:00 and 11:00 results in the lamp being on at 11:00, so that (4) is true, but (5) is just false.

So the original story does not guarantee (4) and (5). One can try to argue that if an infinite number of causes working together is possible, they should be able to work together in a way that produces a Thomson's Lamp for which both counterfactuals *are* true, but this seems contentious.

Alternately, one could argue for an Aristotelian account of laws of nature as grounded in the powers of things. On an Aristotelian picture, (4) and (5) do seem quite plausible given a Thomson's Lamp story. The above argument invoking a law of nature to undercut the two counterfactuals requires a non-Aristotelian law, one not grounded in the powers of things, since we have specified in the Thomson's Lamp story that the only causes are the button presses. Still the argument for (4) and (5) even on an Aristotelian picture may depend too much on an unsatisfactorily weak

intuition that shifts of presses shouldn't change the result but removals of presses should.

2.7 Evaluation

The Thomson's Lamp paradox does not involve any real contradiction on its own. It's only if one adds some further hypotheses, such as the Principle of Sufficient Reason or the counterfactuals (4) and (5), that one runs into difficulties. And even these difficulties seem to rely on an Aristotelian account of causation.

Causal finitism has no difficulty ruling the Thomson's Lamp story out of court, but given the further assumptions needed to make the story be actually paradoxical, the amount of evidence that the story provides for causal finitism is small.

Moreover, as will generally be the case in these paradoxes, the skeptic about causal finitism can also simply block the argument for causal finitism by denying the conditional (1) that if causal finitism is false, then the paradoxical story can be run. However, the conditional *is* intuitively very plausible—at least without further conditions being added to the Lamp story such as (4) and (5). If an infinite number of causes can work together, there seems to be no reason why they couldn't be arranged in the alternating manner indicated by the paradox—though in Chapter 7, Section 3.5 we will consider an alternative way of ruling out the paradox (and the subsequent one) by invoking the discreteness of time, the idea that a finite interval of time (say, an hour) can only contain finitely many times.

Still, the story of the Lamp provides *some* evidence for causal finitism. For instance, there is some reason to accept both the PSR and the Aristotelian picture of laws.

3. Grim Reapers

3.1 Introduction

Thomson's Lamp involved a supertask whose events bunched up at the upper end of the temporal interval. We could reverse this and suppose an infinite sequence of switch toggles at . . . , 10:03.75, 10:07.5, 10:15, and 10:30, with the lamp being off prior to 10:00. Once again, we could ask whether the lamp is on or off at 11:00.

This version of Thomson's Lamp does not have any advantage over the original. Just as the original can (in the absence of an Aristotelian account of laws) be resolved simply by supposing an arbitrary law of nature that determines the outcome or by denying the Principle of Sufficient Reason, the same is true here.

The Grim Reaper paradox, however, puts a twist on the story. In our non-violent version of the paradox,⁶ we suppose the lamp is off at 10:00, and that nothing can (or at least does⁷) turn it on except the pressing of the switch by a Grim Reaper, and

⁶ In some earlier versions, the Reaper had the task of killing someone.

⁷ Theistic readers may worry that necessarily God can turn any lamp in the world on.

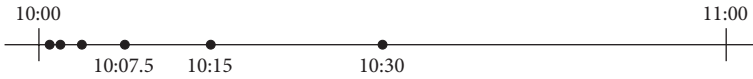


Fig. 3.1 Some representative Grim Reaper activations.

nothing at all can (or at least does) turn it off once it's on. A Grim Reaper is a machine that has an alarm set for a particular time. At that time, the Grim Reaper activates and looks at the lamp. If the lamp is on, it does nothing and goes back to sleep. If the lamp is off, however, the Reaper instantly activates it.⁸

Suppose now that there are infinitely many Grim Reapers (i.e., lamplighters), with activation times at . . . , 10:03.75, 10:07.5, 10:15, and 10:30 (Fig. 3.1).

Then the lamp is on at 11:00. For if it were off at 11:00, then it must have been off at 10:30. But if it were off at 10:30, then the 10:30 Reaper would have turned it on, and so it would have been on at 11:00.

But if the lamp is on at 11:00 and off at 10:00, then how did it turn on? Only a Reaper can turn the lamp on. So a Reaper did. But which one? Let's say it was the one that activated at 10:07.5. Then the lamp must have been off at 10:03.75. But then the Reaper that activated at 10:03.75 would have turned the light on. So the 10:07.5 Reaper couldn't have been the one to turn on the light. And exactly the same argument works for all the other Reapers. None of them turned on the light, but one of them must have, which is a contradiction.

Note that there is something more going on here than in the Thomson's Lamp story. In Thomson's Lamp, until we enriched the story with the PSR or some counterfactuals, all we had was that we couldn't tell from the story whether the lamp was going to be on or off. In the Grim Reaper story, a contradiction follows whether or not the lamp is on or off at 11:00.

3.2 Causal finitism

The causal finitist can rule out the paradox as follows. Each Reaper's activity of deciding (on the basis of observation of the lamp's current state) whether to turn the lamp on is causally prior to the final state of the lamp at 11:00.⁹ Thus, there are infinitely many things impacting causally on one target state, contrary to causal finitism.

And except for worries about infinity, it should be possible to arrange the Reapers as described. Each Reaper is possible, after all, and surely each could go off at any time given between 10:00 and 11:00. We can thus run an argument of a familiar form:

- (8) If causal finitism is false, then the Grim Reaper story is possible.

⁸ If instantaneous action worries readers for reasons of Special Relativity, we leave it to the readers to craft their favorite stories involving particles similar to that in Section 2.4. See Koons (2014) for one such story.

⁹ The details of how the Grim Reapers impact on the final state will be considered when we try to refine the thesis of causal finitism in Chapter 7, Section 2.3.

- (9) But the story is impossible.
- (10) So causal finitism is true.

3.3 *The absurd conclusion objection*

Before going on to consider resolutions and variants of the paradox, let us take the dialectic one step further. One might question (8) precisely because of the success of the argument showing absurdity to follow from the Grim Reaper. The story is impossible, it is contended, simply because an impossibility follows from it (cf. Shackel 2005). In particular, the story would be impossible even if causal infinitism were true.

The absurdity objection is a very powerful response to the argument for causal finitism, and we will have to worry about variants throughout the book.

In response, we can highlight reasons to think (8) is true. Note that the absurdity of the conclusion drawn from the Grim Reaper story is very sensitive to the arrangement of the activation times in the time interval. There is no absurdity in the *Reversed Grim Reaper* story where the Reapers activate at 10:30, 10:45, 10:52.5, and so on (Fig. 3.2). The 10:30 Reaper turns on the light, the later Reapers do nothing, and all is well. And adding a single 10:00 (or 9:59, if one has worries about instantaneous action) Reaper to the original story to get the *Prefixed Grim Reaper* story makes the absurdity of the story disappear: the ten o'clock Reaper turns on the light and the later ones do nothing.

So if the reason why the Grim Reaper story was impossible is *just* the particular absurd conclusion of the story, the Reversed and Prefixed Grim Reaper stories should be possible. Now imagine the herd of Grim Reapers around 9:30 am, and suppose each one has a dial that can be set to adjust its activation time. Then (a) it's impossible to set the dials to the values in the original Grim Reaper story, but (b) any finite bunch of dials can be set to the values in that story, and (c) one can set *all* the dials as in the Reversed and Prefixed stories. This seems wrong. Physical objects like dials should generally be able to be shuffled, moved around, and recombined in minor ways. It is a mark against a theory that it makes the Grim Reaper story impossible but the Reversed and Prefixed ones possible.

Here is another argument that if the Reversed Grim Reaper story is possible, so should the original Grim Reaper story be. For given the Reversed story, in the absence of finitist qualms, we should be able to additionally suppose an infinite number of tinkerers with indeterministic free will¹⁰ adjusting the dials on the Grim Reapers around 9:30. They might all choose to leave the dials alone. But surely it



Fig. 3.2 Some representative reversed Grim Reaper activations.

¹⁰ If one thinks with Hume that freedom requires determinism, then just call it “quasi-free will”.

would be *possible* for them to all set the dials to the settings in the original story. For *each* individual tinkerer could set the dial on *her* Grim Reaper to the setting that it would need to have in the original story. But since the tinkerers are independent and indeterministically free, what other tinkerers are doing doesn't affect what one of them can do. So there should be no difficulty about them *all* setting their Grim Reapers to the values needed for the original paradox. Otherwise, we have to suppose some strange metaphysical force preventing some settings.

One could try to say that the Reversed Grim Reaper story is possible, but only if the dials are not available for infinitely many tinkerers to independently tinker with. But this is implausible.

And the analogue for the Prefixed Grim Reaper story is even harder to defend. To turn a Prefixed Grim Reaper story into a full-blown Grim Reaper story, all we need is a disable button on the 10:00 Grim Reaper and a tinkerer—or even just a random process—capable of pressing that single button. That a Prefixed Grim Reaper story is possible, but that it would be impossible to add the disable button and tinkerer to the story is completely implausible.

Causal finitism, on the other hand, destroys all the variant stories, for the same reason as the original: the final state of the lamp depends on infinitely many events. We could thus put the argument for (8) as follows. If causal finitism is false, then the Reversed and Prefixed Grim Reaper stories are possible. But if the Reversed or Prefixed Grim Reaper story is possible, so is the Grim Reaper story.

3.4 *A rearrangement objection*

Suppose that eternalism is true, so that future events and objects are fully real. Then, as long as finitism is false, it should be possible to have an infinite number of causes, even given causal finitism. After all, a universe that continues forever, with an infinite forwards causal sequence, should be possible. For instance, there is no difficulty about a Grim Reaper activating at noon of each day in an infinite future. Given this, why can't we rearrange the unproblematic activation times of Grim Reapers into the paradoxical ones?

Admittedly, the issue is trickier than the activation of rearrangement times we considered in Section 3.3. There, we imagined that the Grim Reapers exist prior to 10:00 am, and had dials where unproblematic combinations of activation times could be set. That particular thought experiment perhaps cannot be generalized here. The mere existence of infinitely many adjustable Grim Reapers before 10:00 am seems to make for an infinite causal history for the lamp's state at 11:00 am, even if all the activation times are set for after 11:00 am. For that each of the infinitely many dials is set for after 11:00 am causally contributes to the lamp's being off at 11:00 am.

Instead, suppose an unparadoxical story where a Grim Reaper comes into existence at each day of an infinite future, set to go off at noon on that day. Rearrangement suggests that if this is possible, then it should be possible to have all the Grim Reapers come into existence prior to 10:00 am and have their dials set as in the original paradox.

In Chapter 1, Section 3.2, I argued that rearrangement principles should be taken to be defeasible and the best candidates for defeaters are metaphysical principles. Causal finitism is such a principle, and the invocation of rearrangement in the argument from a future sequence of daily Grim Reapers thus has a reasonable defeater.

Furthermore, notice that the rearrangement is not just a straightforward matter of changing the dials. It is also a matter of changing when the Grim Reapers come into existence. In general, temporal rearrangement is intuitively less likely to preserve possibility than spatial rearrangement, the movement of dials, etc. For instance, some versions of Kalām cosmological arguments start with defending the assumption that the past must be finite. Whatever the merits of these arguments, to argue that the past could be infinite because the future could be infinite and a future-directed infinity of events could be rearranged into a past-directed one just does not appear very compelling. Thus, the cost to a causal finitist's denial of this particular instance of rearrangement is lower than the cost in the case of a number of other instances of rearrangement.

The same point goes for a variant (compatible with the denial of eternalism) where we suppose a multiverse with a countable infinity of causally isolated island universes, each containing one Grim Reaper, with the n th world's Grim Reaper set for $60/2^n$ minutes after 10:00 am today. An unrestricted rearrangement principle would allow one to move all these Reapers into a single paradoxical world. But rearrangements that change causally isolated objects into things within a single causal nexus seem more problematic, and should be defeated more easily, than ones that simply rearrange things within a single causal nexus.

It is worth noting that the full finitist is better off with regard to rearrangement arguments: she denies all actual infinities, and it is even more difficult to rearrange the finite into the infinite. But as I argued in Chapter 1, there are good reasons to reject finitism.

3.5 *The mereological objection*

3.5.1 FUSION

Hawthorne (2000) has suggested that in the original story the fusion—the mereological sum or aggregate object made up of all the Grim Reapers—has an effect that no one of the Reapers does. In our context, we would say that the fusion turns on the light, even though no Reaper does.

But when a Reaper sees that the light is already on, it does nothing. We may suppose that it doesn't even touch the switch. The argument we gave implies that the light is on at each of the activation times. So none of the Reapers does anything. And yet mysteriously the light must turn on as a resultant of their joint activity.

There is, admittedly, nothing absurd about a fusion having an effect that no fundamental part does. No particle in a rock breaks the window, but the rock does. However, when the rock breaks the window, each particle in it makes its own tiny contribution, and due to how large Avogadro's number is, these contributions make

for a significant impact. In the case of the Reapers, however, it is stipulated that the individual Reapers do literally nothing beyond observing.

Fusions are supposed to be nothing but the sums of their parts, and so the fusion shouldn't have any causal powers other than those derivative from the parts. (Perhaps Hawthorne's objection can be changed to work with a different kind of whole, an organic whole. On a number of theories of organic wholes—say, Merricks (2001)—an organic whole can have causal powers that go qualitatively beyond the causal powers of the parts. We will consider this option in Section 3.5.2.)

There is a response, however, that invokes Lewis's (1973) counterfactual theory of causation. Take the Grim Reapers out of the Grim Reaper world, so we have a normal world where the lamp is off at 10:00 and stays off. Then we have the counterfactual:

- (11) If we had Grim Reapers activating at . . . , 10:07.5, 10:10, 10:15, and 10:30, the lamp would have been on at all times after 10:00.

On Lewis's theory, such counterfactual dependence between non-overlapping events *is* causal dependence. This yields an argument that the event reported by the consequent causally depends on that reported by the antecedent. But the event reported by the antecedent seems to be precisely the fusion of Grim Reaper activation events.

We should, however, take this to be a *reductio ad absurdum* of the conjunction of Lewis's theory with the thesis that the Grim Reaper story is possible. For it is clear that the Grim Reapers can neither individually nor corporately cause the lamp to be on. It might, of course, be the case that *if* the Grim Reaper story were to hold, the individually immobile Reapers *would* cause the lamp to be on (a little more in favor of this conditional will be said in Section 3.6.3). But since the consequent is impossible, this is reason to deny the possibility of the antecedent.

Here is another way to make vivid the counterintuitiveness of the Hawthorne story. Suppose that the lamp actually has two buttons, a red one making the lamp light up red and a green one making it light up green. Say that a Grim Reaper is, respectively, even or odd provided that it activates at $60/2^n$ minutes after 10:00 for, respectively, an even or odd n . And we now suppose that the even Grim Reapers are programmed only to press the red button while the odd ones are programmed only to press the green one. And both kinds of Grim Reapers do nothing if the lamp is already on, regardless of color.

In the two-color version of the story, much as before, the lamp is on at all times after 10:00, and it is always on in the same color. But which color is that? Both are equally compatible with the story. Now each Reaper is a deterministic cause: its behavior is fully specified. So if the fusion of the Reapers caused the lamp to turn on, then we have one of two amazing things. Either, a fusion of deterministic causes is an indeterministic cause, or else the fusion causes the light to turn on, but we are guaranteed the occurrence of an uncaused event—namely, either the event of the light's being red or the event of the light's being green. Positing uncaused events already counts against a theory (see Section 3.6.2), but being able to

guarantee the occurrence of an uncaused event by doing something (say, activating the Grim Reapers) is even stranger.

3.5.2 NECESSARY EMERGENCE OF ORGANIC WHOLES

Finally, one might also hold that when you get an infinite number of Grim Reapers arranged as in the paradox, you necessarily get a further organic whole that turns on the light, and that's why the original story—which had no such organic unity in it—is impossible.

But there are two kinds of stories in the metaphysics literature about how wholes might *necessarily* come about from parts. The first story is about fusions. It is held by some that necessarily either every plurality of things has a fusion (mereological universalism) while others more weakly think that this necessarily happens whenever one has a plurality of things satisfying some further condition, such as being in mutual physical contact. But these are fusions, and fusions are nothing but the sums of the parts. They are supposed to be a free lunch metaphysically speaking. And what is metaphysically a free lunch does nothing more than the parts do, and is not an organic whole.

The second story is that whenever a bunch of things have the right kind of mutual interdependence or exhibit the right kind of homeostasis, they form an organic whole (1995 talks of the things having a “life” together). This story is quite plausible, but inapplicable to the case at hand. Apart from their commonality of target, the Reapers do not have the right kind of interdependence to form a homeostatic whole.

So it is not plausible that an organic whole would necessarily emerge given the arrangement of Reapers.

3.6 *Uncaused lighting*

3.6.1 OBJECTION

What if one simply says that although the light must be on at all times after 10:00 am, there is nothing that causes it to be on? This commits one to the possibility of contingent events that have a beginning but no cause, contrary to the Causal Principle that events that have a beginning have a cause, i.e., that nothing can come from nothing. Notwithstanding the intuitive plausibility of the Causal Principle, it has been widely denied since Hume (1779) suggested that we can imagine objects coming into existence out of nothing.

My response will have two parts. First, I will briefly review some familiar reasons for accepting the Causal Principle (there is much more discussion of some of the points and objections in Pruss 2006; also, a further argument for a Causal Principle based on a variant of the Grim Reaper paradox will be given in Chapter 9, Section 2.3). Second, I will consider how plausible it is to specifically take the lamp-lighting event in the paradox to be an exception to the Causal Principle.

3.6.2 THE CAUSAL PRINCIPLE IS TRUE

The Causal Principle is very plausible in itself. We shouldn't deny such plausible principles without very good reason.

As a preface, note that certainly we can formulate a narrowly logically coherent first-order logic sentence that says that an object x (for simplicity I will consider events to be a kind of object in this section) came into existence out of nothing:

$$\exists t(\forall u(u < t \rightarrow \sim \text{ExistsAt}(x, u)) \ \& \ \exists u(u \geq t \ \& \ \text{ExistsAt}(x, u))) \ \& \ \sim \exists y(\text{Causes}(y, x))$$

But I argued in Chapter 1, Section 3 that such narrowly logical coherence is not the same as possibility, though of course it is a necessary condition for possibility. Similarly, we couldn't affirmatively settle the question whether something can be self-caused by noting that $\exists x(\text{Causes}(x, x))$ is a narrowly logically coherent sentence of first-order logic.

Two main reasons have been proposed for denying the Causal Principle. First, the fact that we can imagine a situation is a good reason to think the situation is possible. But it is claimed we can imagine a brick coming into existence out of nothing. After all, it seems not hard to imagine: there is nothing and then there is a brick.

Actually, it's pretty hard to imagine that. For it is really hard to imagine nothing. If people are asked to imagine being in a room with nothing (besides themselves) in it, what they imagine is a room empty of the noteworthy features of rooms, like chairs and computers. They don't usually imagine themselves gasping for air and dying in vacuum. Nor do they imagine floating in the room due to the absence of a gravitational field. And quite likely they don't imagine the room as utterly dark, due to the absence of light.

I can, of course, put more effort into my imagining, and try to imagine a dark and airless room without a gravitational field. I have some doubt that I've actually succeeded in this imaginative exercise. Can I really imagine the absence of a gravitational field? Gravitational fields, and *a fortiori* their absences, do not seem to be fit subjects for the imagination. (One can of course imagine something that *represents* a gravitational field, maybe a weird wavy fog.) I suspect that at best what I am imagining is just a dark room that I am mentally labeling as airless and void of gravitational fields. And if I go to the further effort of imagining a truly empty room—empty not just of air, gravitational fields, and light, but also of invisible mythological beings of all sorts, of hypothetical physical fields, and so on—I suspect I am simply strengthening my mental label for the room to “There is nothing *at all* here.”

But of course the possibility of mentally labeling an imagined room as empty gives very little evidence for the possibility of an empty room. One could likewise label a mental image of a person with the tag “This is a married bachelor.”

To imagine an object coming into existence without a cause, however, is even harder than to imagine an empty room. For not only must one imagine the absence of a cause in the locality of the object, one must imagine that there are no distant causes

producing the object via causation at a distance, and no non-physical causes either. This is going far beyond the competence of imagination.

I suspect that in the end the imagination argument comes down to a simple invocation of an intuition: it seems possible for an object to come into existence out of nothing. I do not deny that such intuitions do give evidence. But the intuition is countered by the very widely shared intuition that nothing can come from nothing.

A second line of thought is that science gives us reason to think not only that things *can* happen with no cause at all, but also that they *do* happen with no cause at all.

There are two subarguments here. The first argument is based on particle–antiparticle pairs briefly arising as quantum fluctuations. However, this should not be considered to be a case of something coming from nothing. For the quantum vacuum state from which the pair arises is *something* with a propensity—which can be probabilistically characterized—for the production of such particle pairs. The actualization of this propensity sounds like a causal process.

The second subargument is based on the idea that modern cosmology (and the other arguments for causal finitism, we might add) gives us reason to think that the past is finite. When one combines this with philosophical arguments against an eternal (either infinitely old or timeless) cause of the universe such as God, one gets an argument that something came into existence without cause, namely the universe.

This argument is not very compelling either. To defend the argument one would not only need to argue that God doesn't exist—itself a daunting task, though the problem of evil is at least available as an argument (though see Dougherty and Pruss 2014 for one of many responses)—but one would also have to argue against non-theistic hypotheses about other eternal causes that the universe might have come from.¹¹

3.6.3 IS THE LAMP LIGHTING REALLY UNCAUSED?

The objection to my argument for causal finitism was that the lamp lighting is uncaused. But is the lamp lighting actually uncaused? Clearly had there been no Grim Reapers there, the lamp would not have been lit. On Lewis's (1973) account of causation, the existence of such a counterfactual relation is sufficient for causation. Given that theory, the existence of the Grim Reapers causes the lamp to be on. Similarly, on manipulationist theories of causation (e.g., Woodward 2003), the possibility of manipulating something by means of a correlation is sufficient for causation. But by setting the dials of adjustable Grim Reapers to the paradoxical times versus setting them all to some time after 11:00 am we can control whether the light will be on at 11:00. Thus, on a manipulationist account of causation, the settings on the Grim Reapers do in fact cause the light to be on.

Even apart from such theories of causation, it is plausible that counterfactuals between non-overlapping physical states of affairs are going to hold either due to

¹¹ It is tempting to invoke quantum indeterminism as a third argument for uncaused events. But undetermined quantum events are still *caused* by the physical system in which they happen.

a law of nature that makes one state eventuate when the other does, or due to a causal relation between the states. But we can specify that the Grim Reaper world doesn't contain any additional laws of nature besides those involved in individual Grim Reaper activity. And so it is plausible that we would be pushed to a causal relationship. But it is clear that the lamp's being on doesn't cause the Grim Reapers to exist or activate, so either the Reapers would be causing the lamp to light or there would be a common cause, and in either case the lamp's being lit is caused.

3.6.4 A MYSTERIOUS CORRELATION

Here's another way to think about this. In worlds as in the original paradox, according to the objection at hand, the lamp is causelessly on right after 10:00. In normal worlds where the Grim Reaper activation times are set for a later hour, the lamp stays off past 10:00. Why is there this mysterious correlation so that setting the dials on the Grim Reapers to the paradoxical times somehow non-causally "makes" the lamp go on? Rather than positing such mysterious connections between distinct existences (to echo Hume), isn't it preferable to accept causal finitism? For it will be no surprise that we get strange results when we consider what would happen in the Grim Reaper case if the Grim Reaper case is impossible. And according to causal finitism, the Grim Reaper story—regardless of how the dials are set—is impossible.

3.7 *Discrete time*

Like Thomson's Lamp, the Grim Reaper paradox can also be ruled out if time has to be discrete. We will discuss the discreteness of time in greater detail in Chapter 8. At this point, however, I want to explore a speculative suggestion that the Grim Reaper paradox can be made to run in a discrete setting.

Suppose that the lamp has existed for an infinite amount of time, and that one Reaper activated today, another yesterday, another the day before, and so on. Then it seems we have the same structure as in the original story, but the Reapers are spaced in time in such a way that time can still be discrete.

But now consider how the argument that this is a paradox would go in this scenario. The lamp has always been on, since if it failed to be on n days ago, then it would have been turned on $n + 1$ days ago. And likewise no one of the Grim Reapers activated the lamp.

Where, however, is the paradox? The original paradox was that the lamp came to be on but nothing turned it to be on. However, in the case of an eternal lamp, its being eternally on without anything turning it on seems much less problematic.

But now imagine a scenario on which there are no Grim Reapers, but there is an uncaused lamp that has always been off. Then in that scenario we still have the following counterfactual:

- (12) If there were a Grim Reaper activating on each past day, the lamp would have always been on.

And for exactly the reasons discussed in Section 3.6.3, there is good reason to think that we could elaborate the Grim Reaper story in such a way that this counterfactual dependence of the lamp's state on the Reapers would be causal. But of course the dependence can't be causal, for the reasons discussed in Section 3.5. So there is *something* paradoxical going on given the assumption that the lamp has always been on.

The problem is particularly sharp if even eternal contingent states of affairs—such as the lamp's being on—need to have a cause. Our Grim Reapers are lamplighters. But we can also suppose an infinite sequence of lampdarkeners. Suppose now that some cause *c* caused the lamp to have always been on. Had there been an infinite sequence of lampdarkeners instead of an infinite sequence of lamplighters, then something would have had to cause the lamp to have always been off. But how would the lampdarkeners have prevented *c* from causing the lamp to be on and forced something to cause it to be off?¹²

3.8 Evaluation

The Grim Reaper paradox improves significantly on Thomson's Lamp. It gives us a good reason to believe a hypothesis like causal finitism that allows us to rule the story out of court. Probably the strongest objection to the argument is the objection that the absurdity of the paradox is itself the reason why the story should be ruled out. However, this fails to rule out variants like the Reversed Grim Reaper story, and we have good reason to think that if the Reversed Grim Reaper story is possible, so is the Grim Reaper story. Causal finitism, however, rules out not just the Grim Reaper, but also such permutations. We will, however, have to consider in Chapter 7 whether some other hypothesis may not do a better job than causal finitism, and we will have to refine causal finitism there.

4. Infinite Newtonian Universes

4.1 An argument against causal finitism and a riposte

Here is a quick argument against causal finitism. Newtonian physics is possible. It is clearly possible to have an infinite universe, with infinitely many marbles. But in a Newtonian universe, each object with mass exerts a gravitational force on every other object. The acceleration is proportional to the sum of the forces. Thus every object's acceleration is caused by infinitely many gravitational influences, contrary to causal finitism.

¹² I am supposing here that if there needs to be a cause of the lamp's being on, there needs to be a cause of its being off. But perhaps only positive states of affairs need to have causes. In that case, we can suppose that the lamp has a switch toggling between two colors, red and blue, and we have an infinite sequence of reddenners versus an infinite sequence of bluers, in place of the lamplighters and lampdarkeners.

Surprisingly, this argument can be turned into an argument *for* causal finitism. If causal infinitism is true, the above scenario is possible. But if the above scenario is possible, it should also be possible to have the massive objects be arranged in any other geometrically coherent way, while maintaining Newtonian physics. Here is one such arrangement. Imagine space split in half by an infinite plane (“the central plane”), and call the halves the “left” and “right” halves. Imagine that there is one marble sitting on the central plane, the left half of the universe is empty,¹³ and the right half has an arrangement of masses that has approximately uniform large-scale density. Thus, all large spherical regions in the right half of the space contain approximately the same ratio of mass to volume.

Given this scenario, the marble will experience a gravitational pull from all the masses on the right. This pull will be *infinite*,¹⁴ and hence will result in an infinite acceleration of the marble to the right. Suppose that the other masses are kept roughly in place by other forces, and that there is a narrow corridor for the marble to move through to the right. Then the marble will move to the right, but no matter what finite distance it moves to the right, it will continue to experience an infinite force to the right. Granted, once it moves to the right, it will experience a leftward force from masses now on its left. But that force will be finite.¹⁵

As a result the marble will have to continually infinitely accelerate, and that is impossible—continual infinite acceleration would result in an infinite movement, but where would the marble be after it moved an infinite distance to the right? So the scenario leads to absurdity. Yet if causal infinitism is true, the scenario is possible. So causal finitism is instead true.

Thus Newtonian considerations not only do no damage to causal finitism, but support it. The impossibility of a Newtonian universe uniformly filled on the right with an infinite amount of mass, together with a plausible application of rearrangement, gives us reason to think that any scenario with infinitely many gravitationally interacting masses is impossible, and this in turn gives us a reason to believe causal finitism.

¹³ Or nearly empty, with masses spread out more and more the further one gets from the central plane, if a completely empty half of space is impossible, say because space is relational.

¹⁴ *Gravitational force is inversely proportional to the *square* of the distance, so the pull from further away masses will dissipate. However, the total mass contained within a sphere in the right half of space will be proportional to the *cube* of the sphere’s radius. Imagine then a sphere of space in the right half of space with the marble on its edge, with the sphere having a large radius R . By Isaac Newton’s shell theorem (Schmid 2012), given exactly uniform density, the gravitational force will be the same as the force from an equal mass concentrated in the center. Given approximately uniform density, as in our case, the conclusion will still be approximately true. Since the total mass is $(4/3)\pi R^3 \rho$ where ρ is the density, and the distance to our marble is R , the rightward force on the marble will be greater than or equal to approximately $Gm(4/3)\pi R^3 \rho / R^2 = (4/3)G\pi mR\rho$, where m is the mass of the marble. But just to the right of the marble there are spheres of arbitrarily high radius R with approximate density ρ . It follows that the rightward force is infinite.

¹⁵ We can approximate the masses to the left of the marble by a finite collection of infinite flat planes, and the gravitational force from an infinite flat plane with finite areal density is finite (Hofmann-Wellenhoff and Moritz 2006, p. 135).

4.2 Smullyan's rod

For another Newtonian argument, imagine an infinite flat and perfectly impenetrable terrain, with gravity pulling down uniformly everywhere.¹⁶ Lying on the ground is a perfectly rigid rod with one end right where you are, and that has no other end—the rod stretches to infinity. The terrain is impenetrable. Suppose the rod has a mass of one kilogram. This is quite possible as long as its density drops off with distance from the start: the first meter of the rod can weigh half a kilogram, the second a quarter, the third an eighth and so on.¹⁷

Paradox now ensues when we add that you place your finger under the rod, near the rod's end, and attempt to raise the rod by then gently lifting your finger. If you had a rod that weighed a kilogram and was a centimeter long, you might be able to raise it that way with some good balancing work. But if the rod were a meter long and you tried to lift it up by raising the tip with one finger, you'd fail. The far end of the rod would stay on the ground, and the near end would go up, with the rod swiveling and sliding off your finger. However, if the rod is infinitely long, then amazingly you will be able to do the feat. For since the rod is infinitely long, it will be impossible for it to swivel down from your finger: any downward angle, no matter how small, would cause it to penetrate the impenetrable terrain (Fig. 3.3).

Since you can apply a force sufficient to lift a kilogram with a finger, you will be able to lift the infinitely long rod. It will lift up, perfectly parallel to the ground. That is paradoxical. To see the paradox more clearly, note that as long as the density doesn't drop off too quickly in the first meter or so, at any point in the lift a large unbalanced torque—perhaps even an infinite torque (this is compatible with finite total mass¹⁸)—attempts to swivel the long side of the rod to the ground. Yet the rod stays parallel to the ground.

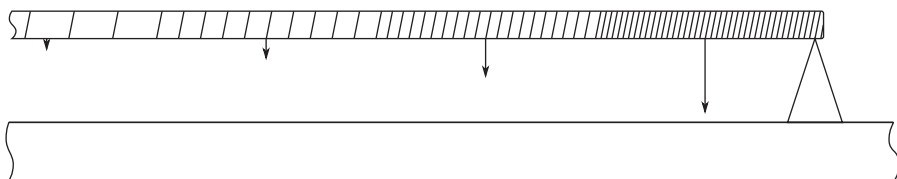


Fig. 3.3 Smullyan's rod with exponentially decreasing density and hence exponentially decreasing quasi-gravitational pull.

¹⁶ Recall that the gravitational pull from an infinite flat plane of uniform finite areal density is finite (Hofmann-Wellenhoff and Moritz 2006, p. 135).

¹⁷ This is a modified version of Smullyan's rod (Smullyan 2008) without requiring infinite strength for lifting it.

¹⁸ *If the density of the rod at distance x is proportional to $1/(x+1)^2$, the total mass will be finite, but the torque about the end will be proportional to $\int_0^\infty x dx/(x+1)^2 = \infty$. On the other hand, if the density at distance x is proportional to 2^{-x} , the total torque will be finite.

But the torque on the pivot is caused by infinitely many instances of gravitational force: the gravitational force on the first meter of the rod, the gravitational force on the second meter, and so on. This violates causal finitism. In fact, more generally, causal finitism makes it impossible to have an infinite rod made of infinitely many parts each of which is subject to a force. Moreover, it looks like the inability of the infinitely many pieces of the rod to penetrate the ground cause the rod to remain stationary (but see the discussion of Benardete's Boards in Chapter 7, Section 2.3 for some complications).

There is a worry that this last argument proves too much, because of a Zenonian extension. Take a rod of length one meter and hold it by one end. Then the torque about the end is caused by the gravitational force on the first half meter, the next quarter meter, the next eighth meter, and so on. Hence it seems that causal finitism rules out even a finite rod. But that's too quick. A finite rod in the real world would be made of a finite number of particles, and so there would be only finitely many gravitational forces to contend with.

But what if we had a one-meter long rod that wasn't made of particles? Classical physics, after all, talked of rigid *continuous* rods. Yet the above Zenonian argument about gravitational forces plus causal finitism still seems to rule out such a rod. And yet classical physics seems to be *possible*.

One answer is this. In the finite case, we shouldn't think of the total torque about the end as caused by an infinite number of actual components. For the sake of computation, we break rods up conceptually into infinitely many parts. But in fact a continuous rod would not have parts—it would be an “extended simple”—or else it would be subdivided in an Aristotelian way. The Aristotelian subdivision would make the rod *actually* have only finitely many real parts, but with each part *capable* of being further split in two. Perhaps the boundaries of real parts are constituted by discontinuities in an extended thing, and so an Aristotelian causal finitist could insist that any rigid rod would exhibit only finitely many discontinuities.

This leads to a further problem, however. What if our original infinite rod is one of those continuous partless rigid objects? (Note: even if it has only finitely many parts, then assuming the parts are contiguous, one of the parts will be infinite, so we can work with the simpler case where the whole is infinite.) In that case, causal finitism would not seem to apply, for the same reason that I argued it wouldn't apply to the finite partless rod.

But there might be additional reasons to be suspicious of *infinite* partless objects in a classical physical setting. To make sense of how finite partless objects interact classically with external forces, we may need global features such as centers of mass. But if a one-directionally infinite partless rigid rod interacting classically is possible, then a bidirectionally-infinite (i.e., extending to infinity in both directions, with no ends) partless rigid rod of uniform density interacting classically should also be possible. But an isotropic bidirectionally-infinite rod has no meaningful

center of mass. All points in it are on par. If causal finitism requires a separate rejection of infinite partless objects in a classical physics setting, that does not seem a great cost.¹⁹

4.3 *The conditional*

The difficulty in many arguments in this book lies in the transition from one possibility claim to another. There are unparadoxical arrangements of infinite amounts of matter in Newtonian space, say ones where the arrangement of matter is less and less dense the further away one goes from some fixed point, and the decrease in density is sufficient to ensure the total gravitational force is everywhere finite. The above argument for causal finitism then requires us to say that if these unparadoxical Newtonian arrangements are possible, the paradoxical arrangement (uniform distribution in one half of space) is also possible.

However, I do not need to defend the conditional if I only want to parry the argument that the possibility of Newtonian physics implies causal infinitism. For our Newtonian paradoxes show that the principle that every mathematically coherent arrangement of mass can be endowed with Newtonian laws is false. But absent such a principle it is open to the causal finitist to say that it is metaphysically impossible for arrangements with infinitely many masses to be endowed with Newtonian laws.

5. Another Eternal Life

If after every day of life there is another day of life, you have eternal life. And if you have eternal life, your life couldn't go on further than it will. But if causal finitism is false, it could: after your eternal life, you could have another eternal life.

The easiest way to see this is with supertasks. Your functioning is sped up by a factor of two, and so you have the first day of your eternal life in twelve hours of external time. Then it's sped up by another factor of two, so you have the second day of your eternal life in six hours of external time. And so on. As a result, you have lived your eternal life—with each day of internal time followed by another day of internal time—all in 24 hours of external time. But then you can go on to have another eternal life after those 24 hours.

What does causal finitism have to do with this? Well, plausibly, our lives are necessarily causally interconnected in such a way that our being alive on earlier days causally contributes to our being alive on later days—this causal interconnection is bound up with the life being a life of a single individual. If, however, you had two eternal lives, one after the other, then your existence during the second eternal life would have among its causes your existence during each of the days of the first eternal life. But that would require causal infinitism.

¹⁹ I am particularly grateful to Ian Slorach for comments on multiple versions of this section.

One may also be able to run this argument without any speeding up. We can give a mathematically coherent description of a time sequence that includes two eternal lives. We just suppose that the temporal dimension is modeled by two copies of an ordinary timeline, with every point of the second timeline coming after the first. Marking the members of the second copy with asterisks, this looks like:

$$\dots, -3, -2, -1, 0, 1, 2, 3, \dots; \dots, -3^*, -2^*, -1^*, 0^*, 1^*, 2^*, 3^*, \dots$$

Now we suppose someone who has a life including the times 0, 1, 2, 3, ... (as well as all intermediate times), and then all the times of the asterisked time sequence. Such a person has an eternal life—and then another. But again this requires causal infinitism if each day of the second life is to be causally connected with all the days of the first.

Such a doubled timeline is very strange. But it is difficult to see why one could rule out the possibility of it, apart from some finitist or causal finitist considerations about the events in it.

6. Time Travel and Causal Loops

6.1 *Grandfathers and togglers*

There is a similarity between the paradox-based arguments for causal finitism and some arguments based on the Grandfather paradox against time travel. The Grandfather argument says that if you could travel back in time, then it would be possible that you go back and kill your grandfather before he met your (corresponding) grandmother, but that would be absurd, so backwards time travel is impossible.

Both the causal finitism arguments and the time travel arguments are subject to the same difficulty, the absurd conclusion objection (see Section 3.3 above). In the time travel case, the objection says that instead of taking the absurdity of killing your grandfather to be evidence against the possibility of time travel, one should simply deny the conditional:

- (13) If time travel is possible, then it's possible that you kill your grandfather before he meets your grandmother

on the grounds that such a killing is absurd.²⁰

I have argued against the absurd conclusion response by showing that there are variants of the Grim Reaper story that have the following two properties: (a) they are not innately paradoxical in the way the original Grim Reaper story is, so the absurd conclusion response fails for them, and (b) if they are possible, so is the Grim Reaper

²⁰ Both the specific version of the Grandfather argument that I am using and this response differ from the argument and response given in the classic paper by Lewis (1976). Lewis considers an argument based on powers—what it would be possible for you to do—and his response is that what it is possible for you to do depends on the contextual background. The present argument is not about powers but metaphysical possibility.

story. The same move is possible in the Grandfather case, and indeed I suspect is a part of the prevalence of the intuition that the Grandfather story is an argument against time travel.

To make the analogy clearer, let's consider a variant of the Grandfather story that is made similar to the Grim Reaper paradox. We have our old friend the lamp with the toggle switch. This lamp suddenly shows up in a dark room at 10:00 am. At 11:00 am a time machine sends it back to 10:00 am, so that it shows up in the aforementioned dark room. The lamp thus embodies a causal loop. So far we have no paradox unless causal loops *per se* count as paradoxes (recall the discussion of regresses in Chapter 2 and note that if an infinite regress is as such absurd, then surely *a fortiori* so is a causal loop). But now imagine a mischievous toggler, who like a Grim Reaper has an activation time, but the togglers, instead of checking the state of the lamp, simply always toggle the lamp's state upon activation. In the paradoxical version of the story, there is only one toggler whose activation time is 10:30 am. Further, we suppose that nothing other than a toggler can affect the state of the lamp—in particular, time travel doesn't affect the state.

We now have a paradox. At 10:31 am the lamp has the same state as the one the toggler put it in. This state doesn't change between then and 11:00 am. Then the lamp travels back so it has the same state at 10:00 am. Nothing affects it, then, until the toggler gets to work. So the state of the lamp before the toggler gets to work is the same as the state of the lamp after the toggler has worked on it, which is absurd.

Again, we can argue that if time travel is possible, the Mischievous Toggler story is possible. But the story is impossible, so time travel is impossible. Once more, the absurd conclusion objection shows up as a denial of the claim that if time travel is possible, so is the Mischievous Toggler story.

There is now a straightforward analogy of our Reversed and Prefixed stories. We can suppose that if a toggler activates when the lamp isn't around—namely before 10:00 or after 11:00—it does nothing. We can add to the Mischievous Toggler story a large supply of extra togglers set to times outside that interval. And now we have a Double Toggler variant, where there are two togglers set to activate at different times between ten and eleven o'clock, say one at 10:30 and one at 10:45, and a bunch of additional togglers set to times outside that interval. There is nothing contradictory about the Double Toggler story. The lamp is in one state at 10:00, then another state from 10:30 to 10:45, and then back in the original state until 11:00.

But once we accept the Double Toggler variant, we really should accept the original Single Toggler story. For it would be really mysterious if the dial on the second toggler had to be set to within the interval between ten and eleven o'clock if the first toggler had an activation time between ten and eleven o'clock but none of the others did. Just imagine the togglers having indeterministic free will, and deciding independently. Suppose all but the last two have decided to set their activation times to outside the ten to eleven interval. We shouldn't suppose that a mysterious force requires that the remaining two either both set their dials to outside the critical interval or both set

them to within. So we should instead think that, just as the Single Toggler story is impossible, so is the Double Toggler story.

This is good reason to deny the possibility of a lamp whose life is a causal loop, and more generally to deny the possibility of time travel and backwards causation of a sort that leads to causal loops.

6.2 *Time travel and backwards causation without causal loops*

The Mischievous Toggler story can be ruled out simply by supposing that causal loops are impossible. If causation by omission would count as part of a causal loop, the same will be true for the original Grandfather story. For then you couldn't travel back in time to a location of spacetime where you would have the power to kill your grandfather prior to his meeting your grandmother, since even if you didn't exercise that power—as of course you didn't, since you really were conceived—your refraining from killing him is causally prior to your existing, but your existing is causally prior to your refraining from killing him.

Should we, thus, take the paradox to rule out all time travel and backwards causation or only the sort of time travel and backwards causation that involves causal loops? The question parallels the question whether the Grim Reaper story should lead us to rule out all infinities or just those that work together causally. In Chapter 1, I argued that we should not rule out all infinities, because doing so creates difficulties in philosophy of mathematics.

I do not have a similarly powerful philosophical argument to limit the restriction on time travel to causal loop cases, though there may be theological arguments based on such cases as prophecy (see Pruss 2007) and praying about past events.

7. Evaluation

The most powerful argument for causal finitism in this chapter is the Grim Reaper argument. It is more convincing than Thomson's Lamp which requires controversial auxiliary hypotheses to make for real paradox. There are also supporting arguments based on Newtonian constructions, as well as based on intuitions about eternal life. Causal finitism (at least when appropriately formulated—this will be refined in Chapter 7) resolves them all, by killing the paradoxes (in the terminology of Chapter 1, Section 1). That is good reason to believe causal finitism.

4

Paradoxical Lotteries

1. Introduction

Imagine a lottery with infinitely many tickets numbered $1, 2, 3, \dots$, and suppose the lottery is fair in the sense that all tickets are equally likely to win, with no ticket privileged over any other.

After reviewing some background information, we will see that such lotteries lead to a number of fascinating paradoxes which we will discuss. The question of how one might go about holding such a lottery is difficult—it's hard to imagine, for instance, picking a ticket uniformly out of an infinite urn in a fair way. But I will offer constructions that show that if causal infinitism is true, then a countably infinite fair lottery is possible. The paradoxes will give us reason to think that a countably infinite fair lottery is not possible, and hence that causal infinitism is not true. Along the way, we will also discuss paradoxical lotteries that are not fair. The strength of the arguments in this chapter depends on how plausible one finds the idea that our modes of probabilistic reasoning should work in infinitary cases.

2. Countably Infinite Fair Lotteries

2.1 Background

Observe that in a fair lottery with more than n tickets, the probability of any particular ticket winning is less than $1/n$. Thus, in a fair lottery with infinitely many tickets, the probability of any particular ticket, say 842, winning is going to be less than $1/n$ for every positive natural number n , since the lottery has more than any finite number n of tickets.

Probabilities, however, are never negative numbers: they must be between 0 and 1. So the probability of getting a particular ticket is a value p such that $0 \leq p < 1/n$ for every $n > 0$. The only real number satisfying this condition is zero itself.¹

Say that a *positive infinitesimal* is a value that is bigger than zero and yet smaller than $1/n$ for every positive natural number n (equivalently, it's bigger than zero and

¹ For if p were other than zero, then we would have $1/p > n$ for every natural number n . Hence $1/p$ would be bigger than every natural number. But there is no real number bigger than every natural number (this is known as the Archimedean property of real numbers): such a “number” would need to be infinite.

smaller than every positive real number). We can then define a negative infinitesimal α as a value such that $-\alpha$ is a positive infinitesimal, and an infinitesimal as anything that is a positive or negative infinitesimal.

Infinitesimals are not real numbers. Nonetheless, it is possible to extend the real numbers in a mathematically rigorous way to add infinitesimals in a way that preserves all the standard arithmetical facts. Such extended systems include the hyperreals (Robinson 1996), the surreals (Knuth 1974), and formal infinite Laurent series (Mendelson 2008, p. 219).² Note that when we add infinitesimals to a system of arithmetic, we also need to add infinitely large numbers. For if α is a positive infinitesimal, so that $\alpha < 1/n$ for all positive integers n , then $1/\alpha > n$ for all positive integers n , so $1/\alpha$ is infinite.

If we allow for the possibility of infinitesimal probabilities, we may not wish to say that the probability of a particular ticket in an infinite fair lottery is zero. Rather, we might say that the probability is zero or infinitesimal. I will henceforth say that x is nearly equal to y provided that $x - y$ is either zero or infinitesimal, and that a proposition or event is nearly certain provided that its probability is nearly equal to 1. It is thus nearly certain that one will lose in an infinite fair lottery if one holds a single ticket—or any finite number of tickets, for that matter.

There is some intuitive reason to prefer the infinitesimal view of the probability of winning in a countably infinite fair lottery to the zero probability view. For instance, it is intuitively more likely that one of tickets 1 and 2 will win than that ticket 1 will win. On the zero probability view, both events have zero probability, while on the infinitesimal view we can say that the probability that 1 or 2 will win is 2α which is bigger than the infinitesimal probability α that 1 will win.

Now it is time to move on to some paradoxes of countably infinite fair lotteries.

2.2 *Expected surprise*

Suppose you enter a fair lottery with a million tickets, numbered from one to a million, and you get ticket number two. That's a surprisingly small ticket number. Most people get a more "representatively sized" number, like 712718 (the number I got when I asked random.org to give me a number from one to a million).

Now suppose you enter our infinite fair lottery, and you get ticket number two. That's surprising (infinitely more!). But if ticket number two is surprising, then what ticket would be unsurprising? Would it be unsurprising to get 712718 or even 10^{712718} ? Such numbers are still incredibly small on the scale from one to infinity. In fact, *any* number you get will be surprisingly small. After all, getting a number as small as 10^{712718} would be incredibly surprising in a lottery with, say, $10^{10^{712718}}$ tickets. But our infinite lottery has infinitely many more tickets than that, so indeed 10^{712718} is extremely surprising. And the same argument generalizes. *Any* number you get will

² *Constructions of the hyperreals and surreals use versions of the Axiom of Choice. Formal infinite Laurent series do not require it.

be surprisingly small in the sense that it was incredibly unlikely that you would get a number that small.

But you can run this argument ahead of time, and so you should rationally expect to get a number that will surprise you by its smallness. That's paradoxical (Hansen n.d.).

It is not paradoxical to expect to be surprised when you do not know what feature of an event will surprise you. Every year, you will see a number of things in the news that will surprise you, and so you can expect to be surprised. But something is or will be rationally wrong if you expect to be surprised *in a very particular way* on a *particular occasion*, say expecting to be surprised at noon next Thursday by getting an Olympic medal. In the infinite lottery case, however, you expect to be surprised in a very particular way: you expect to be surprised by how small the number you get is. Yet nothing goes rationally wrong.

This is not a very compelling paradox. The cost to saying that in infinitary contexts our emotions of surprise will behave oddly is fairly low. It is far from clear, after all, that surprise is the sort of thing that should be governed by rationality. But let us move on to paradoxes that deal more directly with rationality.

2.3 A guessing game

A positive integer N is picked by means of an infinite fair lottery, and it's not revealed to you. Now you are forced to play the following game for each positive integer n : You guess whether N is greater than n ; if you guess correctly, you get a dollar; otherwise, you lose a dollar. Since you have to play infinitely many games, this will take eternity, or at least be a supertask.

For any particular n , clearly you should guess: "Yes, $N > n$." After all, the probability that $N \leq n$ is nearly zero for any finite n . But it is certain that if you follow this strategy for every n , then you will win exactly $N - 1$ times (namely for $n = 1, 2, \dots, N - 1$), and you will lose infinitely many times. That's bad.

In other words, an infinite fair lottery would yield a very simple infinite Dutch Book against a rational agent: a series of gambles each of which is clearly rational to take but which together give a sure loss. This is paradoxical.

2.4 Symmetry

2.4.1 SYMMETRY AND LOTTERIES

Suppose you and I each receive a ticket from a countably infinite fair lottery. There are two ways of filling out this supposition. First, if one can pick a single ticket from a countably infinite fair lottery, one can pick two (e.g., you pick one and run the lottery again with only the remaining tickets). In this case, it's guaranteed that our ticket numbers are different. Second, one could independently run two separate countably infinite fair lotteries. In that case, it's possible that our ticket numbers will be the same,

but it's nearly certain that they won't be. It won't matter for any of my arguments below which option one chooses.

In any case, I then look at my ticket, with surprise noting how small the number is. I realize that there are infinitely many larger numbers and only finitely many smaller (or equal) ones. So I become nearly certain that your ticket number is bigger. But you are in the same boat: you are surprised by how small your number is and are nearly certain that mine is bigger. And we can both predict ahead of time that we will each have such probabilities. There is something very paradoxical about this.³

One way to highlight the paradox is this. Let p be the proposition that my number is smaller and let q be the proposition that your number is smaller. Then you are in as good an epistemic position with respect to q as I am in respect of p . Yet p and q are logically incompatible. Intuitively, neither of us is more likely to be right than the other. When I realize that neither of us is more likely to be right than the other, I shouldn't assign a greater probability to p than $1/2$. Yet I should be nearly certain of p , since once I know the number I got, I am nearly certain that yours is bigger.

Or consider this oddity. You are going to receive a sequence of a hundred tickets from a countably infinite fair lottery. When you get the first ticket, you will be nearly certain that the next ticket will have a bigger number. When you get the second, you will be nearly certain that the third will be bigger than it. And so on. Thus, throughout the sequence you will be nearly certain that the next ticket will be bigger.

But surely at some point you will be wrong. After all, it's incredibly unlikely that a hundred tickets from a lottery will be sorted in ascending order. To make the point clear, suppose that the way the sequence of tickets is picked is as follows. First, a hundred tickets are picked via a countably infinite fair lottery, either the same lottery, in which case they are guaranteed to be different, or independent lotteries, in which case they are nearly certain to be all different. Then the hundred tickets are shuffled, and you're given them one by one. Nonetheless, the above argument is unaffected by the shuffling, since the shuffling does not affect the fairness of the choices: at each point you will be nearly certain that the next ticket you get will have a bigger number, there being only finitely many options for that to fail and infinitely many for it to succeed, and with all the options being equally likely.

Yet if you take a hundred numbers and shuffle them, it's extremely unlikely that they will be in ascending order. So you will be nearly certain of something, and yet very likely wrong in a number of the cases, indeed in about half of the cases. And even while you are nearly certain of it, you will be able to go through this argument, see that in many of the judgments that the next number is bigger you will be wrong, and yet this won't affect your near certainty that the next number is bigger. This is quite paradoxical.

³ This is a variant due to Bartha (2011) of the paradox underlying Freiling's (1986) argument against the continuum hypothesis.

2.4.2 *SYMMETRY AND EXPECTED UTILITY

The reason for raising the above symmetry paradoxes for countably infinite fair lotteries is in aid of an argument that countably infinite fair lotteries are impossible. Later in this chapter, I will argue that if causal infinitism is true, we should be able to have countably infinite fair lotteries, so causal infinitism is to be rejected.

But there is a very similar symmetry paradox that involves no infinite causal histories, and hence is unaffected by causal finitism. Chalmers (2002) offers this two-envelope version of the St. Petersburg paradox:

I am presented with two envelopes, A and B. I am told that each of them contains an amount determined by the following procedure, performed separately for each envelope: a coin was flipped until it came up heads, and if it came up heads on the n th trial, 2^n is put into the envelope. This procedure is performed separately for each envelope. I am given envelope A, and offered the options of keeping A or switching to B. What should I do?

The paradox of course ensues when we note that upon opening envelope A, you find some finite amount of money. But the expected value in envelope B is $(1/2^1) \cdot 2^1 + (1/2^2) \cdot 2^2 + (1/2^3) \cdot 2^3 + \dots = \infty$, which beats whatever finite amount you found in envelope A. So no matter what you find, you should swap. But now we get a serious paradox when we note that you could do the reasoning ahead of time: you know that whatever you find in envelope A, it will be a good idea to swap, and so you might as well swap before finding out what is in envelope A, which is truly absurd. (To add to the absurdity, if you should forget what amount was in envelope A, you would have reason to swap back.)

Chalmers's solution to the paradox is to grant that once you find some finite amount of money, you have reason to swap, but to deny that it follows that you should swap before finding out what is in envelope A. For, Chalmers notes, the reasoning here is dominance-based: no matter what the amount in envelope A is, the expected value of envelope B is better than that. And dominance reasoning is not always right.

But the two-envelope St. Petersburg paradox cuts more deeply than Chalmers realizes. It leads to an invidious two-person Dutch Book. Suppose the house gives you Chalmers's envelope A while your best friend gets envelope B, both for free. You both open the envelopes in separate rooms. The house now offers each of you this independent deal: If you pay a dollar plus *double* the amount in your envelope to the house, the house will give you whatever amount the other envelope contains. Since one dollar plus double the amount in your envelope is still a finite amount, the infinite expected value of the other envelope's amount beats it. So you will go for the deal and your best friend will go for it, too. But a moment's thought shows that the house then ends up handing you and your friend together twice the amount in each envelope, while you and your friend together hand to the house twice the amount in each envelope *plus two dollars*. In other words, the collection of transactions leaves the house ahead exactly by two dollars, no matter what went in the envelopes.

In other words, the two-envelope St. Petersburg paradox can be used to generate a collective Dutch Book against you and your friend. It is plausible that when a Dutch Book can be generated, there must be a failure of rationality among the victims. But it is hard to see any failure here. And the failure that Chalmers identifies in the original paradox is not present here. There is no dominance reasoning. You and your friend are not reasoning ahead of time that they should go for the deals. Rather, the reasoning is happening once you and your friend know what is in the envelopes.

But likewise there are no infinite causal histories in either version of the paradox. The envelopes are filled on the basis of a finite number of coin tosses. There is perhaps a bit of a worry about the unlikely (zero probability) scenario where the coin ends up coming up tails forever,⁴ but nothing in the story causally depends on that scenario. So neither Chalmers's denial of dominance nor the denial of causal finitism helps here.

Note, too, that the Chalmers two-envelope paradox and my strengthening of it can both be run in the case of the countably infinite fair lottery: just suppose that the amounts in the envelopes are set by two independent such lotteries. The resulting paradoxes are closely related to the symmetry paradoxes just discussed. The acceptance of causal finitism is of no help in the new paradoxes, and one might hope that the right resolution for the new paradoxes will help resolve the other countable infinite fair lottery symmetry paradoxes as well, thereby damaging the argument for causal finitism.

However, there is a solution to the two-envelope St. Petersburg paradoxes that does not help with the lottery symmetry paradoxes given above. The St. Petersburg paradoxes are generated by the maximization of expected utility. Empirically, actual human reasoners tend to show risk aversion that does not maximize expected utility, and many decision theorists have argued that they are not irrational in doing so (for an excellent recent account of risk aversion, see Buchak 2014). And Weirich (1984) has proposed the rejection of the maximization of expected utility as precisely the solution to the original St. Petersburg paradox.

But it is important to note that the St. Petersburg paradoxes can still be run on an assumption more general than the maximization of expected utility. Specifically, what the paradoxes need is something like this principle:

- (1) For any probability $p > 0$ and any finite utility M , there is a finite utility N such that probability p of N is better than certainty of M .

Given expected utility maximization, we can see that (1) will be true by taking $N > p^{-1}M$. But some alternatives to expected utility maximization will allow (1) to be true. And as long as (1) is true, we can generate a St. Petersburg-like case, together with its two-envelope variants, using apparently unproblematic assumptions. For fix any finite positive utility M_1 . By (1), let M_2 be a finite utility such that probability $1/2^2$

⁴ I am grateful to Robert Koons for raising this caveat.

of M_2 is better than one dollar plus certainty of $2M_1$. By (1) again, let M_3 be a finite utility such that probability $1/2^3$ of M_3 is better than one dollar plus certainty of $2M_2$. And so on.⁵ Observe that $0 < M_1 < M_2 < \dots$. Now flip our coin and let n be the toss number on which we first get heads. Put a gift certificate for utility M_n in the relevant envelope.

Suppose you open an envelope and find M_n there. Then probability $1/2^{n+1}$ of M_{n+1} is better than a dollar plus certainty of $2M_n$, and the other envelope offers probability $1/2^{n+1}$ of at least M_{n+1} *plus* a non-zero probability of other goodies, so swapping envelopes will be a good deal, even if you have to give back M_n and pay in an extra dollar.

To get out of the collective Dutch Book two-envelope St. Petersburg paradox, we thus need to reject (1). And if Dutch Books are indeed a mark of irrationality, we *must* get out of the paradox, and so the structure of rationality must be such as to rationally force rejection of (1).

Fortunately, we have independent reason to reject (1). Suppose that you are sentenced to ten years of the worst torture sadist physicians can think up. Then a benefactor capable of springing you from prison offers you a choice between (a) release from the torture and (b) probability $1/10^{100}$ of some finite utility N . It is very plausible that no matter how large N is,⁶ nonetheless (a) is the better deal.

Rejecting (1) forces one either to say that there is a finite upper bound to utilities⁷ or to reject expected utility maximization—and any other scheme that could lead to St. Petersburg paradoxes. However, rejecting (1) does no damage at all to the kind of paradoxical symmetry reasoning we previously considered in the cases of countably infinite fair lotteries, if only because (1) is about utilities, while the lottery paradoxes were about credences. Thus we still have reason to reject such lotteries, and the argument for causal finitism is intact.

Moreover, the countably infinite fair lottery version of the Dutch Book two-envelope paradox does not make use of (1). For consider. The amount in each envelope is set through such a lottery on the positive integers. You find n in your envelope. You are nearly certain, then, that the other envelope contains at least $2n + 2$. Now it is clearly rational to pay $2n + 1$ in exchange for near certainty of getting at least $2n + 2$ (and if that's not clear, then note that it's also nearly certain that the other envelope contains at least 2^{2n+1} , and surely it's rational to pay $2n + 1$ in exchange for near certainty of that enormous amount). Risks of finite loss that are nearly certain not to occur can surely be neglected. So, you will rationally pay a dollar plus double what

⁵ ** As formulated, the argument uses the Axiom of Dependent Choice, but if all the utilities are numerical we can eliminate the choices involved by letting M_n be 2 plus the infimum of all M such that probability $1/2^n$ of M is better than certainty of $2M_{n-1} + \$1$.

⁶ At least bracketing the case where N is infinite. Pascal (1858, p. 304) thought that any finite price was worth paying for any non-zero chance of eternal union with God.

⁷ I am grateful to a pseudonymous commenter ("entirelyuseless") on my blog for pointing out this possibility.

your envelope holds in exchange for what the other envelope holds, and so will your friend, and the house will be ahead two dollars. While the St. Petersburg paradoxes depend on tiny probabilities of stupendously large utilities having a large value, as per (1), here we have a case where near certainty of a much larger amount beats the certainty of a much smaller amount. So the countably infinite fair lottery version of the Dutch Book two-envelope paradox gives us reason to reject countably infinite fair lotteries, while rejecting (1) is no help.

2.5 Bayesian manipulation

2.5.1 THE PARADOX

Suppose you're perfectly rational and I've just tossed a hundred fair coins without your seeing the result. I claim that if you are completely sure of the accuracy of my statements and I can run an infinite fair lottery and a certain unparadoxical lottery, I can convince you that nearly certainly all the coins landed heads simply by telling you true things. But a perfectly rational agent should not be subject to such complete manipulation of beliefs by a truthful informer. The possibility of such manipulation is absurd, and hence infinite fair lotteries are absurd.

Here is how my manipulation could proceed. Being perfectly rational, you assign an initial probability of $1/2^{100}$ to the thesis that all the coins landed heads. I now inform you that in an hour I will announce a positive integer number to you. The way that I will generate the number will be as follows. If all the coins landed heads, I will generate a positive natural integer by an unparadoxical procedure that has probability $1/2^n$ of producing number n (note that $(1/2) + (1/2^2) + (1/2^3) + \dots = 1$, so the probabilities correctly sum to one).⁸ But if at least one coin landed tails, I will run an infinite fair lottery and use this to generate a number among $1, 2, 3, \dots$. I will announce the number in both cases but in neither case will I announce which method I used to generate it.

So suppose that you hear me announce some number n , say 15101. Let H be the hypothesis that all the coins landed heads. We can now ask what effect your new evidence E_n , namely that the announced number is n , has on your initially very slim credence in H . To do that, we have to see how likely E_n is on the hypothesis H and on the negation of H . If a piece of evidence is much more likely on H than on the negation of H , then the evidence supports H . Moreover, the degree to which the evidence E_n supports H depends on how much more likely the evidence is on H than on the negation of H .

⁸ One way to do this would be via a supertask if those turn out to be possible despite the arguments of Chapter 3: I toss a fair coin and the number I generate is the number of the toss on which I first get heads, while if I never get heads (which has zero probability) I simply announce the number one. The probability that I first got heads on the n th toss then equals the probability of getting first $n - 1$ tails and then one heads, i.e., $(1/2)^{n-1} \cdot (1/2) = 1/2^n$. Another way is to uniformly choose a real number x in the interval $(0, 1)$ and let n be unique positive integer such that $x \in [2^{-n}, 2^{-n+1}]$. (The length of the latter interval is $(1/2)^n$.)

Now if H is true, then the number 15101 was generated by our unparadoxical procedure such that the probability of getting precisely 15101 is $1/2^{15101}$. On the other hand, if H is false, then the number 15101 was generated by the infinite fair lottery. But the probability of getting a particular ticket number in an infinite fair lottery is nearly zero—it's zero or infinitesimal. But although $1/2^{15101}$ is incredibly small, it is still infinitely many times larger than an infinitesimal, not to mention than zero. So we were infinitely more likely to have generated the number 15101 given H than given the negation of H . It then follows from Bayes' Theorem that once we observe E_{15101} , the probability of H will be nearly one.⁹

More precisely, Bayes' Theorem says that:

$$P(H|E) = \frac{P(E|H)}{P(E|H)P(H) + P(E|\sim H)P(\sim H)} P(H),$$

for any hypothesis H and any evidence E , where $P(A|B)$ is the conditional probability of A given B . We can rewrite the right-hand side as:

$$\frac{1}{1 + \frac{P(E|\sim H)}{P(E|H)} \cdot \frac{P(\sim H)}{P(H)}}.$$

Then if the prior probability $P(H)$ of a hypothesis is neither zero nor infinitesimal while the ratio $P(E|\sim H)/P(E|H)$ is zero or infinitesimal, then $P(H|E)$ will be $1/(1 + \alpha)$ where α is zero or infinitesimal, and it is easy to verify that it follows that $P(H|E)$ will be nearly one.

In the case at hand, the prior probability $P(H)$ is $1/2^{100}$, while $P(E_{15101}|\sim H)$ is zero or infinitesimal and $P(E_{15101}|H) = 1/2^{15101}$. Thus the ratio $P(E_{15101}|\sim H)/P(E_{15101}|H)$ is zero or infinitesimal and hence the posterior probability $P(H|E_{15101})$ of H given the evidence E_{15101} is nearly one. This argument works equally well with any number in place of 15101, so no matter what number I announce, your credence that all one hundred coins landed heads will end up nearly one.

If the possibility of this sort of epistemic manipulation isn't bad enough, there are unfortunate practical consequences. Suppose you are perfectly rational and I am offering you the following wager. If all the hundred coins are heads, I pay you a dollar, but if any of them is tails you pay me a thousand dollars. It would be a terrible idea for you to accept this wager, and so you refuse it. But now I inform you what number I generated using the above procedure, say 15101. If you're rational, then by the above argument you'll be nearly certain that all the coins landed heads, and so your expected payoff for the game will be:

$$(1 - \beta) \cdot \$1 - \beta \cdot \$1000$$

where β is the zero or infinitesimal probability that you now assign to the hypothesis that at least one coin landed tails. But $\beta \cdot \$1000$ is zero or infinitesimal if β is, and

⁹ This paradox is a variant on Dubins (1975).

$(1 - \beta) \cdot \$1$ is \$1 minus at most an infinitesimal, so the expected payoff is \$1 minus at most an infinitesimal, which is definitely worth it.

But of course on most runs of this scenario there will be an occurrence of tails among the hundred tosses. So if we go through the above scenario repeatedly, you will be losing on average close to a thousand dollars,¹⁰ while acting rationally all along.

Once you realize how this is working, you will have reason to stop up your ears prior to my announcement of the number I generated. As long as I can announce that number loudly enough, however, I have a new con game. For a payment of \$900, I offer to refrain from announcing the number (or, even more conveniently for me, I offer not to even bother picking it out), which will allow you to rationally refuse the wager. You will be well-advised to pay up. For if you don't pay up, instead of losing \$900 per run, you'll be losing close to \$1000.

Thus, in situations like this, it is rational for a perfectly rational agent who has no reason to fear loss of rationality to pay not to receive information in order to make better decisions. And that's absurd. Granted, an agent can rationally refuse information that she expects may render her less than rational (say, information about the gender of a job candidate), and would even be reasonable to pay not to get it (e.g., paying an employee to make resumés gender-blind). And even a perfectly rational agent can have a rational preference not to know certain things, whether for instrumental reasons like avoiding spoilers or for non-instrumental reasons like avoiding knowledge of details of parents' intimate lives. But the case at hand is not like those. The reason for your paying \$900 not to receive information is precisely in order to make a better decision, but there is no danger of becoming irrational as a result of the information one is paying not to get. Indeed, it is precisely because one would *remain* rational, so rational that one would accept a wager that an ordinary human agent would refuse, that one pays in this paradoxical case.

It may seem that the potential gains are too trivial compared to the losses for you to be conned. Should you risk a thousand dollars to gain a dollar? But remember that, given the available information, the probability of losing the thousand dollars is at most infinitesimal. And, indeed, it is rational to accept tiny risks of large losses to gain a dollar. For instance, suppose there are two gas stations with slightly different prices. One gas station is half a block further away, but I would end up paying a dollar less. It seems rational to drive half a block to pay a dollar less. But for each additional block that I drive, there is a tiny (but neither zero nor infinitesimal) additional chance that I will die in a car accident, something much worse than the loss of a thousand dollars. And if inflation makes a dollar seem too trivial to be worth thinking about, just scale up my con: if you win, you win ten dollars, and if you lose, you lose ten thousand. Unless I was in a dreadful hurry, I definitely would drive half a block extra to save *ten* dollars, even though there is a tiny risk of a fatal accident while driving that half block. But a fatal accident is much worse than the loss of ten thousand dollars.

¹⁰ More precisely, $(1 - 2^{-100}) \cdot \$1000 - 2^{-100} \cdot \1 , which is \$1000 minus a tiny fraction of a penny.

2.5.2 *A SWITCHOVER POINT?

Recently, Howson (2014) has proposed an interesting solution to the Dubins (1975) version of the paradox, and it generalizes to our setting. The paradox above assumed that when we learn a piece of evidence E , say that the number picked was 15101, we update our credence for H to $P(H|E)$.

Howson proposes what we might call a kind of epistemic satisficing instead. Let X be the announced number. Choose an arbitrary switchover point k , which is not too small. When X is observed to be less than or equal to k , update the credence for H (the hypothesis that all one hundred coins landed heads) to $P(H|X \leq k)$. Otherwise, update the credence for H to $P(H|X > k)$.

Now, as before, by Bayes' Theorem, $P(H|X \leq k)$ will be nearly one, because the ratio $P(X \leq k|\sim H)/P(X \leq k|H)$ will be nearly zero, since if H is not true, then we will generate the number using the countably infinite fair lottery, and that has zero or infinitesimal probability of producing a number less than or equal to k . So, if $X \leq k$, then much as before we will get strong confirmation of $\sim H$. But if $X > k$, this is intuitively how it should be. After all, given H , we wouldn't expect X to be that small (and that works no matter what k is).

On the other hand, one can show that $P(H|X > k)$ will be within an infinitesimal of

$$\frac{P(H)}{P(H) + 2^k P(\sim H)}.$$

This will be a very small number given that in our example $P(H) = 2^{-100}$.

So, on this approach, when we find that X is not too big (i.e., $\leq k$), we do conclude that we got all heads, as we intuitively should, and when X is bigger, we conclude, again as we should, that tails came up at least once.

Nonetheless, the proposal is implausible.

First, suppose that instead of learning what the value of X is, we merely learn that it is greater than k . Given that for every particular value of X greater than k , our credence for H would snap to $P(H|X > k)$, and given that it is antecedently very plausible that when we learn that $X > k$ we should update our credence for H to $P(H|X > k)$, it seems very plausible that upon learning that $X > k$, we should update our credence to $P(H|X > k)$.

This, however, has an unfortunate consequence. For suppose next we learn *what* the value of X is. At that point, our evidence is just as in our earlier discussion, and so by Howson's rule, our credence in H should be $P(H|X > k)$. Hence, learning the value of X provides us with no relevant information about H , once we already know that $X > k$. But that is mistaken. Suppose, for instance, we learn that $X = k + 1$. Once we have updated on $X > k$, the space of possible values of X is $\{k + 1, k + 2, \dots\}$. If H is true, then X was picked by an unparadoxical lottery that assigned twice as great a probability to each number as to its successor. If H is false, then X was picked by a countably infinite fair lottery. Clearly at this point our observation that $X = k + 1$ fits much better with the hypothesis that H is true than with the hypothesis that H is false, and hence our credence in H should go up.

Consider, too, that once we have updated on $X > k$, our situation intuitively looks rather like the original one, but with a shift. For presumably we not only update H by conditioning on $X > k$, but to keep things consistent we condition everything on $X > k$. Let $Y = X - k$. Let P_1 be the probability assignment after updating on $X > k$. Then, if H is false, Y was picked by a countable infinite fair lottery with values in $\{1, 2, \dots\}$. If H is true, Y was picked by an unparadoxical lottery with values in $\{1, 2, \dots\}$, where $P_1(Y = n) = 2^{-n}$. So we have a situation just like the initial one, but with Y in place of X , and the only difference being that now $P_1(H)$ is of the order of magnitude of 2^{-k} , instead of being 2^{-100} . Hence, we should apply Howson's rule once again. Suppose, then, that what we learn is that $Y = 1$. Then since $1 \leq k$, we will have $Y \leq k$, and hence we will update the credence on H to $P_1(H|Y \leq k)$, which will be nearly 1 by an application of Bayes' Theorem as before. Hence, when we first update on $X > k$, and then learn that $X = k + 1$ (which is equivalent to learning that $Y = 1$), our credence in H goes to nearly 1. But our total relevant evidence at this point is that $X = k + 1$, and Howson's rule calls for a different credence in this case.

Next, suppose that after learning that $X > k$, we learn that $X > k + 1$. Well, if learning that $X > k$ led to the posterior credence $P(H|X > k)$ in H , learning that $X > k + 1$ should lead to the posterior credence $P(H|X > k + 1)$ in H . Surely, something relevant has been learned, and we need to update our credence on it. But Howson cannot say this. For suppose we *first* learn that $X > k + 1$ and our credence in H goes to $P(H|X > k + 1)$. Next suppose we learn that in fact $X = k + 2$. At that point, the equation $X = k + 2$ summarizes our total evidence, and so our credence in H would be given by Howson's rule as $P(H|X > k)$. But $P(H|X > k) > P(H|X > k + 1)$. Hence, the claim that $X = k + 2$ would be evidence against H after having learned that $X > k + 1$. But that would be mistaken. It is the unparadoxical lottery, which occurs only when H is true, that would favor $X = k + 2$ among $\{k + 2, k + 3, \dots\}$.¹¹

2.5.3 *COUNTABLE ADDITIVITY AND CONGLOMERABILITY

In the setting of classical probability theory, Good's Theorem (Good 1967) guarantees that it never pays for a perfectly rational agent who has no reason to fear loss of rationality to refuse free information in order to make better decisions. Our paradoxes, however, do not contradict Good's Theorem, since classical probability theory assumes countable additivity of probabilities, which is violated by countably infinite fair lotteries.

Indeed, the paradoxes we just discussed are fundamentally due to the lack of countable additivity in the lottery probabilities. A probability function P is countably additive provided that whenever E_1, E_2, \dots are disjoint events, then $P(E_1 \vee E_2 \vee \dots) = P(E_1) + P(E_2) + \dots$. Classical mathematical probability theory assumes all probability functions to be countably additive. But in the countably infinite fair lottery, we

¹¹ Howson (2014) also suggests a different method of updating involving smoothing rather than a sharp cut-off. The details of that would need to be worked out before evaluating how that fares with the objections above.

do not have countable additivity. The reason we do not have countable additivity differs depending on whether the probability of a particular ticket winning is zero or infinitesimal.

If the probability is exactly zero, then we lack countable additivity because $1 = P(E_1 \vee E_2 \vee \dots)$ if E_n is the probability of ticket n being picked (it's certain that some ticket or other is picked) whereas $P(E_1) + P(E_2) + \dots = 0 + 0 + \dots = 0$.

If, on the other hand, $P(E_n) = \alpha$ for some (positive) infinitesimal α , then things are more complicated. The standard systems for construction of infinitesimals do not in general define a countably infinite sum of infinitesimals, at least in our case where the summands are the same.¹² Thus, the required equation $P(E_1 \vee E_2 \vee \dots) = P(E_1) + P(E_2) + \dots$ does not hold, since although the left-hand side is defined, the right-hand side is not. In our infinite fair lottery case, we can intuitively see why we shouldn't be able to have a meaningful sum. For consider our infinite sum:

$$\begin{aligned}\alpha + \alpha + \alpha + \alpha + \dots &= (\alpha + \alpha) + (\alpha + \alpha) + \dots \\ &= 2\alpha + 2\alpha + \dots \\ &= 2(\alpha + \alpha + \dots).\end{aligned}$$

If the value of this sum is x , then $x = 2x$. But if x is not zero, then we can divide both sides by x to yield $1 = 2$, and so x must be zero. However, x cannot be zero since it must be at least as big as α , and hence a contradiction follows from the assumption that the sum has a value.¹³

The lack of countable additivity in the case of an infinite lottery is responsible for a phenomenon known as non-conglomerability. A probability function P is conglomerable with respect to a partition E_1, E_2, \dots (a partition is a collection of pairwise disjoint events such that their disjunction is the whole space of possibilities) provided there is no event A and real number a such that for all i we have $P(A|E_i) \leq a$ and yet $P(A) > a$. Conglomerability is a very plausible property. Suppose you are certain that some event in the partition will occur. If you also know for sure that whatever event in that partition you learn occurs, your probability for A will be at most a , then how could your rational probability for A be more than a ?

Conglomerability is closely related to van Fraassen's very plausible Reflection Principle which says that if one is rationally certain that one will have a certain rational credence, one should already have that credence now (van Fraassen 1984).

But typically, where there is no countable additivity, there is lack of conglomerability (Schervish, Seidenfeld, and Kadane 1984). In the case of the countably infinite

¹² *For formal Laurent series, one can always add term-by-term, as long as the term-by-term sums all converge. But they won't converge when the summands are all non-zero and equal.

¹³ Admittedly, we cannot regroup all infinite sums. Specifically, we cannot regroup conditionally convergent series. But intuitively when all the summands are positive—and this is a theorem in the case of real-valued summands—we should be able to regroup to our hearts' content. And the point here is only to feed intuition.

fair lottery, we can see the lack of conglomerability directly. Let E be the event that the ticket picked will be even and O the event that it will be odd. By *finite* additivity, $P(E) + P(O) = 1$, so at least one of the two events must have probability at least $1/2$. (Intuitively, they both have probability exactly $1/2$, but I don't need that for the argument.) Suppose that $P(E) \geq 1/2$ (the argument in the case where $P(O) \geq 1/2$ will be very similar). Then consider the partition provided by the following sets:

$$\begin{aligned} E_1 &= \{2, 1, 3\} \\ E_2 &= \{4, 5, 7\} \\ E_3 &= \{6, 9, 11\} \\ E_4 &= \{8, 13, 15\} \\ &\dots \end{aligned}$$

Observe now that each set E_n contains exactly one even number and two odd ones. Thus, by the fairness of the lottery, $P(E|E_n) = 1/3$. Thus, $P(E|E_n) < 1/2$ for all n , but by assumption $P(E) \geq 1/2$, and conglomerability is violated.

Where conglomerability is absent, one gets strange results such as reasoning to a foregone conclusion and paying not to receive information (Kadane, Schervish, and Seidenfeld 1996), just as we saw in Section 2.5. And the symmetry puzzle in Section 2.4 is also a non-conglomerability puzzle. Taking the original two-ticket version, the probability that my ticket number is bigger than yours is initially within an infinitesimal of $1/2$. But the conditional probability that my ticket number is bigger than yours given what my ticket number is—whatever that may be—is at most an infinitesimal, and so conglomerability is violated.

One possible response to my preceding paradoxes is that non-conglomerability needs to be accepted when dealing with countably infinite fair lotteries, and non-conglomerability just happens to have a number of paradoxical consequences. But the cost of accepting non-conglomerability is high, namely many paradoxical consequences. It is better to take non-conglomerability in these lotteries to be both a paradox in its own right and the mathematical root of a number of other paradoxes.

2.6 Improving everyone's chances

Suppose I have a fair lottery with five tickets and one winner. I can keep within these parameters and improve *most* people's chances of winning. For instance, I can make the first ticket have zero chance of winning, and then raise the other tickets' chances from $1/5$ to $1/4$. But the following principle is very plausible:

- (2) In a lottery where a prize is awarded to exactly one winner, it is impossible to change the probabilities of victory so as to improve *everyone's* chances of winning.

In classical probability, (2) is a theorem (in the special case of countably many players, it is an immediate consequence of countable additivity). Still, once one allows for

infinitesimal probabilities, one might become suspicious of (2). For instance, suppose that a lottery has uncountably many tickets, one ticket corresponding to each real number between 0, inclusive, and 360, exclusive. A spinner is spun, and the winning number is given by the angle in degrees between the initial and final positions of the spinner. Moreover, the spinner is spun in such a way that the winning number is uniformly distributed on $[0, 360)$.

Now suppose that we modify how the winning number is generated as follows. Once the spinner comes to a stop, we double the angle, and then rewrite it so that it's still within $[0, 360)$. For instance, if the angle where the spinner stops is 195 degrees, we double it to give 390 degrees, and then note that 390 degrees represents the same angle as 30 degrees, so our winner is 30 degrees. But now the ticket numbered 30 has *two* ways of winning: it wins if the spinner stops at 15 degrees as well as when the spinner stops at 195 degrees. The same is true for every other ticket. So our modification to the lottery doubled everyone's chances of victory, even though it was still the case that only one ticket wins.

The classical probabilist agrees with what I just said, but says that this is compatible with (2). For the probability of the spinner coming to rest at a particular exact angle is zero, and when one doubles zero, one still has zero, which is no improvement. However, if one thinks that the probability of an outcome is not zero but an infinitesimal, then one must admit that the above procedure has increased everyone's probability of winning. I am inclined to think that (2) is sufficiently plausible that this argument should lead us to the classical conclusion that each number has *zero* probability. But some might instead be led to reject (2).

Given that (2) is highly plausible, someone who rejects it should try to make that rejection more palatable. A reasonable approach is to say that our intuitions are not sensitive to infinitesimal differences in probability. Thus, it may be possible to *infinitesimally* improve the chances of victory, but no more than that. We should thus replace (2) with this weaker claim:

- (3) In a lottery where a prize is awarded to exactly one winner, it is impossible to change the probabilities of victory so that each individual's chance of victory becomes higher by a non-infinitesimal increment.

It is very hard to deny (3). But if countably infinite fair lotteries are possible, then it must be rejected. For suppose we have a countably infinite fair lottery with tickets 1, 2, 3, Now consider a second lottery with the same tickets, but where the probability of ticket number n winning is $1/2^n$. Within classical probability theory, this is an entirely unparadoxical lottery, and easy to construct.¹⁴ But now observe that in the second lottery, every ticket has a non-infinitesimally bigger chance of winning. For ticket n has an infinitesimal chance α of winning on our countably infinite fair

¹⁴ See note 8, above.

lottery, but a non-infinitesimal chance $1/2^n$ on our second lottery, and so the increase in the probabilities of winning is $1/2^n - \alpha > 1/2^n - 1/2^{n+1} = 1/2^{n+1}$, and hence is non-infinitesimal. Hence, countably infinite fair lotteries contradict (3), and thus are impossible.

The paradox of violating (3) is different from the preceding in that it does not appear to be as closely related to non-conglomerability.

3. Constructing Paradoxical Lotteries

3.1 Fairness and paradoxicality

All the paradoxes above arise from one crucial consequence of fairness in a countably infinite fair lottery: for any finite set of tickets, the probability that the winning ticket is a member of that set is zero or infinitesimal (i.e., is nearly zero). This, in turn, follows from the finite additivity of probabilities and the fact that each ticket has at most an infinitesimal probability of winning.

More generally, I will stipulate that a lottery is *paradoxical* provided that there are countably infinitely many tickets and each ticket has at most an infinitesimal probability of winning. The arguments given so far in this chapter justify the use of the word “paradoxical”, and give us good reason to think that paradoxical lotteries are impossible. However, I will now argue that if causal infinitism is true, then it is possible to have a paradoxical lottery. One or two (depending on whether one thinks the construction in the next subsection is a cheat) of the constructions will even yield a *fair* countably infinite lottery.

3.2 Lucky coin-flip sequences

Given causal infinitism, toss an indeterministic fair coin infinitely often, either in a supertask or in an infinite past. Real-life coin tosses may be deterministic, so “coin toss” may need to be a stand-in for some quantum experiment. (I will at times omit the word “indeterministic.”) Number all the flips with the positive integers (e.g., if you did one flip a day over an infinite past, number today’s flip 1, yesterday’s 2, and so on). Here is something that *might* happen: one and only one flip yields heads (Fig. 4.1). If you’re lucky enough that this happened, then let n be the number associated with the flip that landed heads. This number n then is the number chosen by the lottery. And the choice is fair: the tosses are all on par, so no number is privileged. We have what Norton (2018) calls “label independence”.

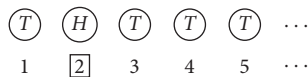


Fig. 4.1 A lucky case where the lottery works, with the winner being the number 2.

Of course, this method for generating a lottery typically fails. The Law of Large Numbers guarantees that with probability one, on average half of the coin-flips will be heads. But in real-life lotteries, it's also possible for the process to fail. The organizer may be about to pick a ticket out of a hat, but a tornado blows the hat away. When we say that a lottery is fair, we mean that conditionally on the lottery being successful, all outcomes are on par. And in *this* sense, when we get lucky and have exactly one heads in the infinite sequence, we really do have a countably infinite fair lottery. The difference is that in real-life cases, failure of the lottery is unlikely, while in this case, success is unlikely.

Nonetheless, while this simple construction really does generate a lottery when it's successful, the unreliability of the construction makes the paradoxes less telling. For instance, the manipulation argument in Section 2.5 becomes less impressive if I have to be extremely lucky to be able to manipulate you, since only if I am extremely lucky am I going to be able to draw on the result of a countably infinite fair lottery. Indeed, perhaps there is some plausibility in thinking that a perfectly rational agent might fail in such rare circumstances.

One might try to make the construction work more often by repeating it until it succeeds. Unfortunately, it turns out that even if we run the construction countably infinitely often, we still cannot expect it to succeed. In classical probability, the probability of success will still be zero (Norton and Pruss 2018), as a countable disjunction of zero probability outcomes has zero probability.

Another way to make use of repetition and luck to generate an infinite lottery machine is as follows. Flip a countably infinite number of coins, but arrange them in a two-dimensional array, with infinitely many rows, each of infinitely many flips. Do not look at the outcomes of the flips. Instead, send a robot to traverse the array in a supertask, in a zigzagging way, say like in Fig. 4.2. The robot determines whether it is the case that every row contains exactly one heads, and informs us of this.

If the robot returns a negative answer—i.e., if some row contains a number of heads other than one—then our attempt to make a lottery machine has failed. But if the robot returns a positive answer, the robot is then directed to the beginning of the first row of the array. The coins lying in the array together with the robot now constitute a lottery machine. To operate it, the robot is placed to the left of the first row, and then

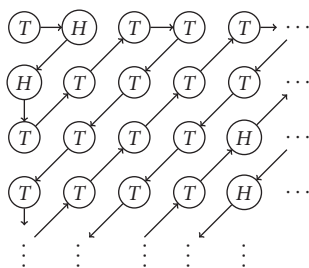


Fig. 4.2 A traverse of a two-dimensional array.

is directed to move to the right until it reaches heads, report the position of the heads, and finally move to the beginning of the next row.

Once the lottery machine is constructed, it can generate as many countably infinite fair lottery tickets as we like, enabling all the paradoxes, including the ones that depend on multiple lotteries. Of course, we had to get lucky enough for each row to contain exactly one heads, and the probability of such luck was zero. But a machine that needs luck to construct is still a constructible machine. And, to be honest, any infinite machine requires something like great luck to construct: one would need to be very lucky to have the material resources for an infinite machine, as well as never to make a mistake in building it.

It is interesting that the existence and operation of the machine constituted by the coin array and the robot does not violate causal finitism—the robot only has to traverse finitely many tails to get to a heads in each row. But in order to generate the paradoxes of countably infinite fair lotteries, we need to *know* that we have such a lottery in hand. And the “quality control” step where the supertasking robot first checked that all the rows of the array conform to the requirements depended on causal infinitism.

In Section 3.4, I will consider a more complex way to generate a countably infinite fair lottery. That more complex method will depend on the possibility of causally exploiting the Axiom of Choice instead of just getting improbably lucky. In Chapter 6, I will argue that if causal infinitism is true, then such causal exploitation is possible. But, first, let us reflect on what it is to construct a countably infinite fair lottery.

3.3 *What it is to construct a countably infinite fair lottery*

Countably infinite fair lotteries violate the axiom of countable additivity within classical probability theory. This makes it challenging to say what it is to construct such a lottery, since one cannot make use of standard probabilistic constructions. I will solve this problem by invoking symmetry.

Intuitively, some stochastic processes, considered as real physical causal processes, are *symmetric*. Thus the space Ω of outcomes of the process has a natural set of symmetries, which are one-to-one functions of Ω onto itself, with the property that if σ is a symmetry and A is a subset of Ω , then the event of the outcome being in A and the event of the outcome being in $\sigma A = \{\sigma(\omega) : \omega \in A\}$ are on par stochastically.

If we have a probability measure P on Ω , then a necessary condition for A and B to be on par stochastically is that $P(A) = P(B)$. But this necessary condition may not be sufficient. For instance, in classical probability, if you are throwing a dart uniformly randomly at a continuous circular target, then the probabilities that the dart will hit the exact center of the target and that the dart will hit the horizontal line going through the exact center are equal: they are both zero. But nonetheless intuitively the two possibilities are not on par: it seems infinitely more likely that the dart will hit the horizontal line through the center than that it will hit exactly the

center. Furthermore, the sets A and B may be on par but non-measurable in classical probability.¹⁵

There are non-classical probability theories that allow one to make such comparisons; for instance, Popper functions and infinitesimal probabilities (these two approaches are basically equivalent, as per Krauss 1968 and McGee 1994), or comparative probabilities (Fine 1973). But they all have their technical shortcomings,¹⁶ and it is better for the notion of the symmetry of a stochastic process and the interrelated notion of sets of outcomes being on par not to be tied to any particular account. Thus, I will simply take the notions of a set of symmetries of the space of outcomes of a stochastic process and of sets of outcomes being on par as not further analyzed, but as interrelated.¹⁷

As an example of symmetries of a causal process, consider a countably infinite sequence of fair, independent, indeterministic coin tosses, all on par physically. We can then take an outcome of this process—i.e., a sequence of heads or tails—and reorder it in some fixed way. For instance, perhaps each odd-numbered toss is swapped with the succeeding even-numbered toss, so that $TTHTTTHH\ldots$ is mapped to $TTTHHTTH\ldots$. Such reordering of the elements in the sequence is intuitively a symmetry of the system, and sets of outcomes that differ by such a reordering are intuitively probabilistically on par. In particular, we would expect that the probability of getting such-and-such a frequency of heads among the even-numbered tosses equals the probability of getting that frequency among the odd-numbered tosses.

A different symmetry that will be particularly useful to us in the case of a sequence of coin tosses is reversal of a set of outcomes. Suppose our coin tosses are numbered in some way. Then given any set S of numbers, and given a sequence ω of coin toss outcomes, we can define $\sigma_S\omega$ as the sequence ω with every outcome of a coin toss whose position in the sequence of tosses is in S being switched to the opposite value, tails to heads and heads to tails. In other words, if $\omega = (\omega_1, \omega_2, \ldots)$ is a sequence of coin tosses and $\sigma_S(\omega) = (\omega'_1, \omega'_2, \ldots)$, then $\omega'_i = \omega_i$ if and only if $i \notin S$. For instance, $\sigma_{\{1,4,5\}}(TTHTTTHH\ldots) = HTHHHHHH\ldots$. Intuitively every σ_S is a symmetry of a sequence of fair and indeterministic coin tosses, because it would make no difference probabilistically speaking if, after generating the coin tosses, we were to go through and mechanically turn over every coin corresponding to a number in S , as long as S was specified without relying on the outcomes of the coin tosses.¹⁸

We can now specify more precisely what I will mean by generating a countably infinite fair lottery out of a stochastic process. We start with a stochastic process with a space of outcomes Ω and a set of symmetries G such that A and σA are always

¹⁵ E.g., imagine a spinner that uniformly chooses a point on a circle. Let A be a non-measurable subset of the circle, and let B be any rotation of A (even A itself): then A and B are on par, but neither has a probability.

¹⁶ For instance, see Pruss (2013a and 2015).

¹⁷ Cf. Bartha and Johns (2001) for a related symmetry approach.

¹⁸ It would obviously be unfair to say that we turn over all and only the coins that are heads, as that would guarantee that we have only tails.

probabilistically on par for a subset A of Ω and $\sigma \in G$. Then we find a function f that assigns to each member ω of Ω a positive integer $f(\omega)$ such that for any positive integers n and m , there is a symmetry σ in G such that $\sigma E_n = E_m$, where $E_i = \{\omega : f(\omega) = i\}$.

Our procedure for running the lottery then is now this: You run the stochastic process, plug the outcome ω of it into f , and announce the positive integer $f(\omega)$ as the winning number. Thus, $f(\omega)$ is the winning number of the lottery at location ω in the space of outcomes, and E_i is the event of ticket i being the winner. Then fairness is ensured by the requirement that for any pair of positive integers n and m , the events E_n and E_m correspond under a symmetry of the space of outcomes and hence are on par.

In the above, I assumed that for *every* outcome of the stochastic process we can generate a winner. But one could also suppose that there is a proper subset U of Ω and that the function f is defined only on U , while keeping all the other conditions. In this case, if the original process lands outside U , no lottery is generated. But if it lands at a point ω within U , we are lucky, and we can take $f(\omega)$ to be the winner of the lottery. In this case, we have an incompletely reliable construction of the lottery.

In Section 3.2, we had a particularly unreliable construction. For simplicity, let us work with the single-row variant. To fit this construction into a symmetry-based framework, we will generate a lottery whose outcomes are arbitrary integers rather than just positive ones, and we suppose that the coin tosses are indexed by all integers, with the obvious extension of our framework. One obvious family of symmetries for the coin toss process is given by shifts τ_n , where n is an integer and $\tau_n(\omega)$ is the sequence ω shifted by n .

Let Ω be the set of all bidirectionally infinite sequences of heads or tails, i.e., all functions from the set of all integers (positive and negative) \mathbb{Z} to $\{H, T\}$, thought of as generated by infinitely many fair and indeterministic coin tosses. Let U be the subset of Ω consisting of all sequences that are heads in all but one position. For $\omega \in U$, let $f(\omega)$ be the position of the unique tails. Then if m and n are any two integers, $\tau_{m-n}\{\omega : f(\omega) = n\} = \{\omega : f(\omega) = m\}$, since ω has its unique tails in the n th position if and only if a shift of ω by $m - n$ has its unique tails in the m th position. Thus we do have a way to construct a countably infinite fair lottery, albeit one very incompletely reliable in that it only works if the original process generates a sequence in U , and that is exceedingly unlikely.

3.4 *Coin-flips and the Axiom of Choice

A more technical construction of an infinite lottery depends on the Axiom of Choice from set theory in place of luck.

Start with an infinite sequence of independent fair coin-flips, either in a supertask or an infinite past, and suppose that the flip numbers are numbered $1, 2, 3, \dots$. We will denote a heads flip with a 1 and a tails flip with a 0, and so we can model the situation with the state space Ω of all sequences $(\omega_1, \omega_2, \dots)$ where each ω_i is zero or

one. A particular run of the coin-flips then generates a particular sequence of zeroes and/or ones.

Let σ_n be a function that takes a sequence $(\omega_1, \omega_2, \dots)$ in Ω and reverses the value of the n th item in the sequence. Thus,

$$\sigma_n((\omega_1, \omega_2, \dots)) = (\omega_1, \omega_2, \dots, \omega_{n-1}, 1 - \omega_n, \omega_{n+1}, \omega_{n+2}, \dots).$$

I will in general assume that if ω denotes a sequence in Ω , then ω_n is the n th element of it.

The functions σ_n are intuitively symmetries of the process that generates the sequence of coin-flips. If n is any positive integer, U any subset of Ω , and $\sigma_n U = \{\sigma_n(\omega) : \omega \in U\}$, then the coin-flip sequence's being a member of U is stochastically on par with the sequence's being a member of $\sigma_n U$.

Let Ω_0 be the subset of Ω consisting of those sequences that have only finitely many ones in them. There is a one-to-one correspondence between Ω_0 and the natural numbers \mathbb{N} . To see this, note that, given $\omega \in \Omega_0$, we will have $\omega = (\omega_1, \dots, \omega_n, 0, 0, \dots)$ for some n . Then let ω^* be the natural number which can be written in binary $\omega_n \omega_{n-1} \dots \omega_1$. For instance, if $\omega = (0, 1, 0, 1, 1, 0, 0, 0, \dots)$, then ω^* is the number whose binary expansion is 11010 (or, equivalently, 011010 or 0011010—it doesn't matter how many trailing zeroes from ω we include), i.e., the decimal number $2 + 8 + 16 = 26$. Hence, we can interpret the members of Ω_0 as natural numbers written backwards in binary.

Write $a \oplus b$ for the sum modulo 2 of two numbers a and b from $\{0, 1\}$. Thus $0 \oplus 0 = 1 \oplus 1 = 0$ and $0 \oplus 1 = 1 \oplus 0 = 1$. If ω and ω' are in Ω , then let $\omega \oplus \omega' = (\omega_1 \oplus \omega'_1, \omega_2 \oplus \omega'_2, \dots)$. For any subset U of Ω and any $\omega \in \Omega$, write $\omega \oplus U = \{\omega \oplus \omega' : \omega' \in U\}$. We can think of $\omega \oplus \omega'$ or $\omega \oplus U$ as a “twist” of ω' or U , respectively, by ω .

If $\omega \in \Omega_0$, then twisting by ω is equivalent to applying a finite number of σ_n transformations, namely those σ_n transformations where n corresponds to a non-zero entry in ω . More precisely, if n_1, n_2, \dots, n_k are the distinct indexes of the non-zero positions in ω , then $\omega \oplus \omega' = \sigma_{n_1}(\sigma_{n_2}(\dots \sigma_{n_k}(\omega')))$. Since combinations of symmetries will be symmetries, the twist operations will be symmetries of our stochastic process.

Now, define an equivalence relation \sim on Ω by saying that $\omega \sim \omega'$ if and only if $\omega \oplus \omega' \in \Omega_0$. This is the same as saying that ω and ω' agree except in at most finitely many places, and it's clear that this is an equivalence relation. Let $[\omega] = \{\omega' : \omega \sim \omega'\}$ be the equivalence class of ω .

By the Axiom of Choice, suppose that g is a function that assigns to each equivalence class A some member $g(A)$ in A . We are now ready to describe a countably infinite fair lottery. For simplicity, the ticket labels will be the members of Ω_0 , which correspond to the natural numbers (and hence the positive integers) by backwards binary encoding. Given a sequence ω of our coin-flip stochastic process, let $f(\omega) = \omega \oplus g([\omega])$. In other words, $f(\omega)$ is a sequence that has a 1 precisely in those places where ω differs from the specially chosen representative $g([\omega])$. Since $\omega \sim g([\omega])$, it

follows that $f(\omega)$ has only finitely many ones in it, and hence $f(\omega) \in \Omega_0$. Our lottery works as follows: an infinite sequence ω of flips is generated in an infinite past or via a supertask, and the winning ticket is $f(\omega)$.

It remains to show that if μ and ν are any two elements of Ω_0 , then $\{\omega : f(\omega) = \mu\}$ and $\{\omega : g(\omega) = \nu\}$ are equivalent under some symmetry. To see this, note that $\nu \oplus \mu \in \Omega_0$, and hence all we need to show is that:

$$(\nu \oplus \mu) \oplus \{\omega : f(\omega) = \mu\} = \{\omega : f(\omega) = \nu\}.$$

To show this, observe that \oplus is commutative, associative, and each member of Ω is its own inverse, i.e., $\alpha \oplus \alpha = \mathbf{0}$, where $\mathbf{0}$ is the sequence that is always zero, and that the map that sends ω to $\omega' = \nu \oplus \mu \oplus \omega$ is a bijection of Ω onto itself that is its own inverse. Observe also that $[\nu \oplus \mu \oplus \omega'] = [\omega']$, since $\nu \oplus \mu \in \Omega_0$. Then:

$$\begin{aligned} (\nu \oplus \mu) \oplus \{\omega : f(\omega) = \mu\} &= (\nu \oplus \mu) \oplus \{\omega : \omega \oplus g([\omega]) = \mu\} \\ &= \{\nu \oplus \mu \oplus \omega : \omega \oplus g([\omega]) = \mu\} \\ &= \{\omega' : \nu \oplus \mu \oplus \omega' \oplus g([\nu \oplus \mu \oplus \omega']) = \mu\} \\ &= \{\omega' : \omega' \oplus g([\nu \oplus \mu \oplus \omega']) = \nu \oplus \mu \oplus \mu\} \\ &= \{\omega' : \omega' \oplus g([\nu \oplus \mu \oplus \omega']) = \nu\} \\ &= \{\omega' : \omega' \oplus g([\omega']) = \nu\} \\ &= \{\omega : \omega \oplus g([\omega]) = \nu\} \\ &= \{\omega : f(\omega) = \nu\}. \end{aligned}$$

Therefore, $\{\omega : f(\omega) = \nu\}$ and $\{\omega : f(\omega) = \mu\}$ are equivalent under a twist by $\nu \oplus \mu$, which is a member of Ω_0 , and hence are equivalent under a symmetry of our original stochastic process.

Thus we indeed have a construction of a countably infinite fair lottery.

3.5 Random walks

A robot moves along an infinite line. At each step in its career, it flips an indeterministic fair coin. On tails it moves one step left and on heads one step right.

Suppose that the robot has gone through an infinite sequence of steps, stretching backwards in a reversed supertask¹⁹ or an infinite past. If we number the positions on the line with integers, the robot's current position can be taken to yield the outcome of a lottery.

The lottery is paradoxical. This is because the probability function for the robot spreads out over time, and over an infinite amount of time it will spread out so as to ensure that the probability of its being in any one position is zero or infinitesimal.

¹⁹ E.g., one step taken at 11:00:00, one at 10:30:00, one at 10:15:30, and so on, as opposed to a forwards supertask where one step is taken at, say, 10:00:00, another at 10:30:00, another at 10:45:30, and so on. Norton (2018) criticizes the forwards supertask version of this story.

Slightly more rigorously, let $p_k(x)$ be the probability that at the k th step the robot is at position x , where $k = 0$ corresponds to the last step, $k = -1$ the previous, and so on. Fix a very large natural number N . Then the probability function p_0 will depend on p_{-2N} . Fix an integer x_0 , and let's ask how large $p_0(x_0)$ can be. The biggest $p_0(x_0)$ can be is when p_{-2N} assigns probability 1 to x_0 and 0 to all others: while it is unlikely that a robot that starts at position x_0 will again be at x_0 after $2N$ of steps, it will be more unlikely yet that a robot that starts anywhere else will be at x_0 after $2N$ steps.²⁰ Now, if $p_{-2N}(x_0) = 1$, and hence all the other probabilities at step $-2N$ are zero, then the only way the robot can be at x_0 again at step 0 is if exactly N heads and N tails were tossed.

Let $f(N)$ be the probability that in $2N$ tosses of a fair coin, exactly N heads are tossed. We have just shown that no matter what p_{-2N} is, $p_0(x_0) \leq f(N)$. And this argument works for every natural number N . But it follows from facts about the binomial distribution that $f(N)$ converges to 0 as N goes to infinity.²¹ Thus for any real number $r > 0$, if N is large enough we will have $f(N) \leq r$, and thus, $p_0(x_0) \leq r$. Hence $p_0(x_0)$ must be zero or infinitesimal. And this is true for every x_0 .

Thus every ticket number has zero or infinitesimal probability, and so we have constructed a paradoxical lottery.

The above scenario assumed that space could be infinite. That assumption is not essential. We could, instead, suppose a particle or other object that has a property—say, spin or charge—that changes randomly upward or downward with equal probabilities and that has no upper or lower bound.

4. Objections

4.1 *Infinite lotteries and uniform distributions*

4.1.1 THE PROBLEM

There is, however, a serious problem in using paradoxical lotteries as an argument for causal finitism. For it seems that one can generate paradoxical lotteries without infinite causal histories. Suppose that we have a process—say, a spinner—that uniformly generates a real number in the interval $[0, 1)$. We can use the outcome of a single run of this process to power the constructions we just gave. The constructions in Sections 3.2 and 3.4 only required a countably infinite sequence of coin tosses. But take our random real number. If the n th digit after the decimal point is even, then deem the n th coin toss to be heads; otherwise, deem the n th coin toss to be tails.

²⁰ *This follows from the fact that the binomial distribution for probability $1/2$ peaks at the mean.

²¹ *We have $f(N) = 2^{-2N} \binom{2N}{N} = 2^{-2N} (2N)! / (N!)^2$. By Stirling's Formula, $n!$ is asymptotic to $\sqrt{2\pi n} (n/e)^n$. Thus, $f(N)$ is asymptotic to $2^{-2N} (2\pi N)^{-1/2} (2N/e)^{2N} / (N/e)^{2N} = (2\pi N)^{-1/2}$, which converges to zero.

Given this infinite sequence of deemed coin tosses, it seems one can generate a lottery exactly as before, but without using any infinite causal sequence.²²

4.1.2 RESPONSE I: NO CONTINUOUS DISTRIBUTIONS

One solution is that in addition to our learning from the paradoxes that infinite causal sequences are impossible, we also learn that it is impossible to have a causal process that uniformly generates a real number in an interval.²³ This would require us to basically reject all continuous distributions, since all other continuous distributions can be used to generate uniform ones.²⁴ In particular, we couldn't have anything like continuous spinners, darts thrown at continuous targets, or radioactive decay processes where the decay time ranges over real numbers.

This is not so crazy. Causal finitism makes it at least plausible that time is discrete (see Chapter 8). A discrete time is either infinitely subdivisible or not. Aristotle thought it was infinitely subdivisible: between any two points of time there might not *actually* be another time, but another time is nonetheless possible. Infinitely subdivisible time might be enough to allow a radioactive decay to generate a continuous distribution. For while the infinitely many times at which an atom might decay don't all exist, perhaps there is a continuum of *possible* times at which it could decay, and if it *were* to decay, that would make one of these possible times actual. But it is not clear that even this would generate a *continuum* of potential decay times as opposed to, say, just a countable number of candidate times.

But it is also a serious possibility that time has to be made up of basic intervals, not further subdivisible. If so, then a continuous-time decay process will be impossible, ruling out decay processes and spinners generating a continuous distribution.

Furthermore, it is plausible that space is like enough to time that if time has to be discrete, so does space, and if time has to be discrete with non-subdivisible intervals, space is similar (see Chapter 7, Section 2.2). And that would rule out dart-throwing processes being used to generate a continuous distribution.

²² There is a minor issue about what to do in the cases of numbers that have two different decimal expansions. For instance, $0.50000 \dots = 0.49999 \dots$. We could always choose the decimal expansion ending in an infinite string of zeroes over one that ends in an infinite string of nines. This will at most infinitesimally affect the relevant probabilities, because the probability that one will hit a decimal number that has two different expansions is zero or infinitesimal, since all such numbers are rational, and the probability that the uniform choice will hit a rational number is at most infinitesimal. I will thus ignore this complication.

²³ This introduces a minor dialectical complication in that I used a spinner to generate an unparadoxical lottery where the probability of ticket n was 2^{-n} . But I could have also used a supertask (toss coins until you get heads and let n be the number of the toss on which you first got heads), and our causal infinitist couldn't very well object. We would still have a *reductio ad absurdum* against causal infinitism.

²⁴ *Suppose a random variable X , within classical probability, has a continuous distribution, so that the function $F(x) = P(X < x)$ is continuous. Consider now the random variable $Y = F(X)$. Consider any $y \in (0, 1)$. Let x be the smallest real number such that $F(x) = y$ (this exists because of the continuity of F and since F has limits 0 and 1 at $-\infty$ and $+\infty$ respectively). Then $Y < y$ if and only if $X < x$. Hence $P(Y < y) = P(X < x) = F(x) = y$, and it follows by countable additivity that Y is uniformly distributed in $[0, 1]$.

4.1.3 RESPONSE II: MEASUREMENT OF INFINITE PRECISION DATA

Moreover, to make use of a continuous distribution to generate an infinite number of coin-flips seems to require an infinite number of measurements. To check if the first digit after the decimal point of our random variable X is even or odd requires measuring X to a precision better than ± 0.05 , and to check if the next digit is even or odd requires a precision better than ± 0.005 , and so on. For simplicity, let's say we are working with the "lucky" construction of Section 3.2. Then we need to verify the equivalent of the claim that exactly one coin landed heads, i.e., we need to verify that exactly one digit of X after the decimal point is even. This seems to require making more and more precise measurements, presumably in a supertask, and then collating the data obtained in these measurements in order to find out whether exactly one digit is even. And here causal finitism will be violated, as all the measurements will be in the causal history of one's knowledge of the collated result.

We could, perhaps, suppose a law of nature or causal propensity of the system that directly ensures some effect, say a bell ringing, whenever every digit of X after the decimal point is even, so that there is a single cause— X being such that exactly one digit after the decimal point is even—of that effect. Or perhaps there could be a divine being who simply announces whether this is indeed the case. This won't work, however, in a quantum mechanical context where the facts about the value of X are caused by the measurements, and where an infinitely precise value may not make any sense—for instance, an infinitely precise position would require a non-normalizable momentum distribution.

So if we are willing to say that all measurements of continuous data have to be akin to quantum mechanical cases, we have a way of showing that paradoxical lotteries generated via continuous random selections would have to involve an infinite number of measurements, and hence would be ruled out by causal finitism.

4.1.4 RESPONSE III: THE USE OF THE AXIOM OF CHOICE

There is yet another move that we can make in the case of the reliable construction of Section 3.4, though admittedly not in the less impressive "lucky" construction of Section 3.2.

The reliable construction of countably infinite fair lotteries out of coin-flips invoked the Axiom of Choice, namely the claim that if S is a set of non-empty sets, there is a function that chooses an element of every element of S .

However, it is crucial to the construction that the constructions are not merely mathematical constructions, but that they exhibit something that, given causal infinitism, could be a genuine causal process. And to that end, it's not enough that there abstractly be a choice *set*. Rather, it is necessary that there be a causal process that implements the choice function. In our constructions, we needed a process that, on inputting a particular sequence of heads and tails, would produce as output the chosen member of the equivalence class containing that input. But it is plausible that

such a process would need to have an infinite complexity—it would need to match the input sequence against the equivalence classes, and then output the right pre-chosen representative of the class. Thus the process would need to have access to a pre-chosen representative for each of the equivalence classes, perhaps stored in an infinite memory store. Given causal infinitism, there will be no difficulties here. But given causal finitism, it is far from clear that this can be done. We will have more discussion of the Axiom of Choice and its implications in Chapter 6.

Furthermore, simply recognizing that two sequences fall in the same equivalence class requires examining infinitely many members of the sequences, which is apt to violate causal finitism.

4.2 *A non-normalizable quantum state

John Norton (2018) has considered a quantum mechanical construction of a countably infinite fair lottery. A quantum mechanical state $|\psi\rangle$ is normally represented as a superposition, i.e., sum (or, more generally, integral), of orthogonal basis states $|\phi_n\rangle$:

$$|\psi\rangle = \sum_{n=1}^{\infty} c_n |\phi_n\rangle$$

where $|\phi_n\rangle$ has norm 1 and $\sum_{n=1}^{\infty} |c_n|^2$ is finite. If we then have an observable O that is certain to have value n in state $|\phi_n\rangle$, then a measurement of O in state $|\psi\rangle$ has probability $|c_n|^2$ of yielding n .

Norton now imagines a system in a non-normalizable state, one where $\sum_{n=1}^{\infty} |c_n|^2 = \infty$. Such a system may not be nomically possible, but nonetheless may be a metaphysically possible physical system, much like a violation of matter-energy conservation would be a physical event that is nomically impossible but metaphysically possible.

If this is right, then it should also be possible to have a system where the coefficients c_n equal 1 for all n . Since the ratio of $|c_n|^2$ to $|c_m|^2$ (when $c_m \neq 0$) gives the ratio of the probability of a measurement of O yielding n to the probability of its yielding m , if all the coefficients are equal the probabilities are all equal, and we have a countably infinite fair lottery.

This construction does not appear to involve an infinite number of causes: the cause of the measurement of O yielding its result is the *single* state $|\psi\rangle$ and the process of measurement. Granted, the state $|\psi\rangle$ can be expressed mathematically as the sum of the infinitely many states $|\phi_n\rangle$, but the latter states do not actually exist, so it does not seem that there are infinitely many things in the causal history. If the above is correct, then we can have paradoxical lotteries without causal infinitism, and the argument that we need causal finitism to block the lotteries fails.

Norton notes, however, that there are multiple technical problems with the non-normalizable state construction. Besides discussing specific problems with physical realizations of the setup, Norton notes that $\sum_{n=1}^{\infty} |\phi_n\rangle$ is not a member of the Hilbert

space of states over which quantum physics is defined. I take the failure of the sum state to be a member of that Hilbert space to be as serious a problem with the construction as it would be to consider a Newtonian scenario where there is a particle at coordinates $(1, 1, 1) + (1, 1, 1) + (1, 1, 1) + \dots$. It's not just nomically forbidden for a particle to be at (∞, ∞, ∞) , but Newtonian physics just has no sense assigned to those coordinates. Likewise, quantum physics assigns no sense to $\sum_{n=1}^{\infty} |\phi_n\rangle$.

Indeed, the mathematics of Hilbert spaces does not assign any sense to this sum. An infinite sum is defined as the limit of finite partial sums, but the partial sums $\sum_{n=1}^N |\phi_n\rangle$ do not have any limit as N goes to infinity. We can, of course, embed the Hilbert space in a larger mathematical space with respect to whose topology the partial sums do converge. But it does not seem that the larger mathematical space has any *physical* meaning, so that specifying that the system is in a state that is in that space may simply be physical nonsense, much as it would be physical nonsense to suppose that a quantum state is a non-vector, like the number 7.²⁵

A physical system needs to be defined in a way that assigns something like chances to different ways for the system to evolve. We could, then, extend our Hilbert space to include non-normalized states like $\sum_{n=1}^{\infty} |\phi_n\rangle$ and then try to add physical meaning to the extension by specifying that the ratios of chances of measurements continue to correspond to ratios of squares of absolute values of coefficients. But there are two problems with this. First, it is not clear that we have specified a physical system when we have merely given ratios of chances rather than the actual chances. Second, the specification simply assumes, without any argument, that it is metaphysically possible to have a system which both (a) includes the state $\sum_{n=1}^{\infty} |\phi_n\rangle$ and (b) is governed by the rule that ratios of chances equal ratios of squares of absolute values of coefficients. And to assume this is essentially to assume that countably infinite fair lotteries are possible. Orthodox Quantum Mechanics does nothing here to make the assumption plausible.

Thus, while the non-normalized state construction is clever, and does not appear to use causal infinitism, it does not give us reason to think that one can construct paradoxical lotteries without violations of causal finitism.

4.3 *Limitations on our reasoning*

Perhaps, though, the difficulties with paradoxical lotteries—as well as the paradoxes of rationality considered in Chapter 5—simply show that our reasoning abilities are incapable of dealing with certain kinds of infinite situations. Maybe there are situations when there is no well-defined probability after conditionalization on the

²⁵ **We can take the Hilbert space H of quantum states and extend it mathematically to the larger non-Hilbert topological space $H \cup \{7\}$ whose topology is generated by the topology of H plus the singleton $\{7\}$. We can even extend the resulting topological space further to be a vector space, by embedding it in some larger space that has the requisite properties of a vector space. But this still gives no meaning to the physical state of the system being the number 7.

evidence, or perhaps conditionalization is not the appropriate way to go. This should not be a surprise: our reasoning abilities evolved to deal with problems in the earth ecosystem that do not involve infinities.

This, however, is a risky line of thought. Modern science heavily uses infinitary mathematics, especially in the guise of differential equations. If our reasoning becomes unreliable when infinities appear on the scene, we cannot trust science, which is an unacceptable conclusion.

Perhaps, though, one could distinguish between deductive reasoning about infinities, as in the case of mathematics, and probabilistic empirical reasoning, as in the paradoxes above. And it could be that as a contingent matter of fact our world is finitary in nature, and infinities need only occur in mathematical idealizations.

However, even if our world is finitary in nature, if this is a merely contingent fact, we need empirical reasons to reject various infinitary hypotheses if these are metaphysically possible. And so we would need to be able to reason probabilistically about infinitary hypotheses in order to reject them. Hence, even if we lived in a finitary universe, we would need to be able to think both deductively and inductively about metaphysically possible infinitary hypotheses.

Furthermore, even apart from the special case of infinities, the line of thought that limits our reasoning to the kinds of scenarios that came up in our evolutionary history threatens to undercut much too much. While it could be that the reasoning in the modern sciences is similar in kind to the kind of reasoning our ancestors engaged in when they strove to escape predators and capture prey, the context in which the reasoning is applied is very different. To trust science, we have to have confidence that our reasoning is not limited to the narrow contexts where it evolved.

Moreover, if an imaginable scenario is metaphysically possible, an agent could consistently think that she is in such a scenario, or at least think that she has a non-zero probability of being in such a scenario, and the puzzles about what to do or think in that case could become puzzles for her.

Finally, the paradox that if a countably infinite fair lottery is possible, then one can increase everyone's chance of winning, and not merely infinitesimally, seems to be more than just a paradox about rationality and epistemic probabilities, but a paradox about objective chances, and hence may not be affected by the limitation-on-reasoning response. Of course, one might take the limitation response sufficiently broadly as to undercut all intuitions about things beyond our experience. But that would lead to skepticism in metaphysics, science, and even mathematics (think of intuitions about the axioms of set theory).

5. Evaluation

Countably infinite fair lotteries are paradoxical in my technical sense that every outcome has infinitesimal or zero probability. There are multiple ways of seeing that a

lottery that is paradoxical in this sense is impossible. It leads to expected surprise, symmetry puzzles, Bayesian manipulation of rational agents, and the possibility of improving everyone's chances at winning. We should thus reject the possibility of paradoxical lotteries. Yet it is very plausible, in light of several different constructions, that if it is possible to have infinite causal processes, it is also possible to have paradoxical lotteries. Hence infinite causal processes are impossible, and we have an argument for causal finitism.

Once again, a challenge to this line of thought could come from the idea that perhaps *these* particular causal processes are impossible, and impossible because they give rise to such paradoxes, but *other* infinitary causal processes are perfectly fine. However, this puts implausible fine-grained limits on what causal systems can possibly be instantiated. And in particular it is *very* implausible that one could have an infinite causal history and yet it would be impossible to have a random-walk process like that in Section 3.5.

The second main challenge is that countably infinite fair lotteries can apparently be generated by innocent-seeming continuous distribution processes apparently without any infinite causal histories. Because of issues involving the Axiom of Choice, this only affects the "lucky" construction of Section 3.2. Moreover, the causal finitist who wants to rely on paradoxical lotteries can point to worries about the possibility of genuinely continuous distributions and infinitely precise measurements.

Perhaps the strongest challenge, however, is the claim that inductive reasoning like ours suffers from significant limitations in infinitary contexts. An appropriately formulated version of this thesis can resolve most of the paradoxes of this chapter. (A possible exception is the paradox that one can non-infinitesimally increase everyone's chance of winning in a lottery.) But causal finitism has an advantage over this solution: it is also supported by metaphysical intuitions about infinite regresses (Chapter 2) and resolves paradoxes that do not involve rationality (Chapter 3).

5

Probability and Decision Theory

1. Introduction

In this chapter we shall consider several paradoxes in probability and decision theory that do not depend on infinite fair (or paradoxical) lotteries, in order to see what support they give to causal finitism.

We begin with some compelling diachronic and synchronic die-guessing paradoxes. For instance, it will turn out that if one has access to information about infinitely many *past* die rolls, and causal infinitism is true, then one should be able to make sure that one usually guesses die rolls correctly. However, an independently interesting “parody” setup where causal infinitism does not appear to be at issue will also need to be evaluated.

Finally, we consider two paradoxes that are less well resolved by causal finitism: Satan’s Apple and Beam’s paradox. These will turn out to have multiple variants. On some variants they will not actually depend on causal infinitism, and hence causal finitism is of no help in resolving them. But these variants will also make the paradoxes easier to resolve in other ways, whereas the variants that appear most compelling can indeed be resolved with causal finitism. The evidence for causal finitism provided by Satan’s Apple and Beam’s paradox is not very significant, but we need to evaluate the paradoxes to ensure that they do not weaken the overall case for causal finitism by artificially limiting the usefulness of causal finitism.

2. Guessing with Finitely Many Errors

2.1 *Doing a little better than one can*

Every day, over an infinite past, a fair, indeterministic die has been rolled, each time independently and with no memory of past rolls, and you have no information about future rolls. The experiment comes to an end on some specific day in the future. Before each time the die is rolled, you are asked to guess whether the die will show a six. If you get the answer right, you got a treat. If you get it wrong, you get an unpleasant electric shock. Moreover, the results of the die rolls are causally independent of your guesses.

No matter how much information you have about past outcomes of die tosses, as long as you are completely sure that the die is fair and each roll is independent, your optimal strategy is:

- (1) ALWAYS NO: Always guess "Not a six".

For guessing "Six" gives you a $5/6$ chance of the shock while guessing "Not a six" gives you only a $1/6$ chance.

But if causal infinitism is true, then you could have access to all the past rolls, and if you have that, then there is a strategy that is better than ALWAYS NO. Say that a number came up "almost always" (or on "almost all" rolls) provided that it came up in all but at most finitely many cases. Then our better strategy is:

- (2) ALMOST ALWAYS NO: If almost always a six came up, guess "Six"; if infinitely many non-sixes came up, guess "Not a six".

Strategy ALMOST ALWAYS NO is better than ALWAYS NO. For it is just as good as ALWAYS NO in those cases where the infinite sequence of rolls contains infinitely many non-sixes. In those cases, both strategies always recommend guessing "Not a six". But in the admittedly unlikely cases where the infinite sequence of rolls contains only finitely many non-sixes, following ALWAYS NO results in getting an electric shock almost always. On the other hand, ALMOST ALWAYS NO results in getting a treat almost always. Thus, for some sequences, ALMOST ALWAYS NO does better than ALWAYS NO and for the others it is equivalent. Hence it is the better strategy of the two.

This yields a quick argument for causal finitism. In our scenario:

- (3) One cannot use past information about the rolls of a fair memoryless die to get a better guessing strategy than the best strategy that does not use past information.
- (4) If causal infinitism is true, one can use past information about the rolls of a fair memoryless die to get a better guessing strategy than the best strategy that does not use past information.
- (5) Therefore, causal infinitism is false.

Premise (3) is a generalization of the denial of the Gambler's Fallacy. The standard Gambler's Fallacy holds that if some result has come up often, it's less likely to come up again. But of course a die doesn't remember that it's come up a certain way, so the Fallacy is a fallacy. We generalize the point to say that past information is of no help when one knows for sure that the dice are memoryless and fair. Premise (4) is true since given causal infinitism, one could have the information needed to make ALMOST ALWAYS NO work (for instance, one could set up a device that lights up if and only if there were infinitely many non-sixes).

Causal finitism, then, tells us that although ALMOST ALWAYS NO may be a superior *mathematical* strategy, it cannot be implemented in practice, and so there is no problem for rationality.

There is another way to see that ALMOST ALWAYS NO is better than ALWAYS NO. No matter what sequence of rolls occurs, ALMOST ALWAYS NO is at least as good. But it also has the advantage of *guaranteeing* that one will be right infinitely often. For if there are infinitely many non-sixes, ALMOST ALWAYS NO, just like ALWAYS NO, ensures you are right infinitely often. But if there are only finitely many non-sixes, ALMOST ALWAYS NO will make one be right almost always, and hence infinitely often, while ALWAYS NO will make one be wrong in all but at most finitely many cases.

In fact, it is in and of itself paradoxical that there be a guessing strategy for a fair and memoryless die that *guarantees* being right infinitely often—or even guarantees ever being right.

One might worry that a really clever mathematician might come along and generate a strategy that guarantees infinitely many correct guesses but that doesn't depend on causal infinitism—a strategy that requires information only about finitely many guesses—and if so then the argument for causal finitism would have little force. But it turns out that there is no such strategy. In fact, there isn't even a strategy based on a finite amount of past information that guarantees that you're *ever* right (though of course, it's very likely that you are). This follows from the following result, once we notice that information on which exact non-six numbers came up is irrelevant¹ and hence we can encode the relevant sequences of tosses as zeroes, for when a non-six was rolled, and as ones, for when a six was rolled.

THEOREM. *For every sequence of natural numbers n_0, n_1, n_2, \dots and functions f_k from n_k -tuples of numbers in $\{0, 1\}$ to $\{0, 1\}$, there is a sequence $\dots, c_{-2}, c_{-1}, c_0$ of numbers in $\{0, 1\}$ such that for all $k \leq 0$ we have $f_k(c_{k-n_k}, \dots, c_{k-1}) \neq c_k$.*

The Theorem says that for any guessing strategy $(f_k)_{k \leq 0}$ for zeroes and ones where at step k we have information about the past n_k results, there is a sequence c_k of zeroes and ones such that the strategy always fails. The proof of the Theorem is given in the Appendix to this chapter.

2.2 A contradiction

ALMOST ALWAYS NO generates a decision-theoretic paradox. But it seems we can also generate an outright contradiction along similar lines. The ALMOST ALWAYS NO strategy has the property of ensuring that you are infinitely often right. This property is not dependent on the die being fair or memoryless, nor on the fact that the die is rolled before you guess. Thus, this property should remain even if someone places the die in front of you after you guess. Suppose, then, that you adopt ALMOST ALWAYS NO, but an enemy who wishes you ill then places the die in front of you in a configuration

¹ More precisely, any strategy s that guarantees at least one correct guess and that makes use of information on which particular non-six numbers came up can be replaced by a strategy s^* that does not make use of this information and that guarantees at least one correct guess. To see this, let s^* be the strategy obtained by taking at each stage the past data, replacing all non-sixes other than 1 with 1, and feeding that into s , and using s to decide whether to guess "Six" or "Not a six". If s is guaranteed to yield at least one correct guess, so is s^* .

different from the one you guessed (six if you guessed non-six, and non-six if you guessed six). Given causal infinitism, this all seems possible: you have a well-defined strategy and so does the enemy.

But this is a contradiction. For ALMOST ALWAYS NO guarantees that you're right infinitely often, while your enemy's simple but effective strategy ensures that you're never right. Yet it is very plausible that if causal infinitism is true, then the ALMOST ALWAYS NO strategy can be implemented no matter what the die placement strategy is. And your enemy's strategy can be implemented no matter what guessing strategy is used against it.

Perhaps, though, here one can respond in the same way some would respond to the Grim Reaper: the story is contradictory, and hence impossible. A dialectic similar to the one in Chapter 3, Section 3.3 can now ensue, regarding the plausibility of rearrangements of strategies that do not lead to contradiction (say, the strategy of the other party—now, a friend—placing the die precisely as one has guessed) into strategies that do lead to contradiction.

I will, however, focus on the decision-theoretic versions of the paradox in this chapter, rather than the contradiction-based version, in order to broaden the number of *types* of paradoxes that lead to causal finitism. The decision-theoretic paradoxes do not involve an outright contradiction, but a violation of the principle (3) of the rational uselessness of past information in cases of independent experiments, and it seems less plausible to reject the setup because of such a violation than because of an outright contradiction.

2.3 *Doing much better than one can*

The ALMOST ALWAYS NO improvement on ALWAYS NO is enough to yield (4), but it is not particularly impressive. After all, it is only in the extremely unlikely scenario that there are only finitely many non-sixes that ALMOST ALWAYS NO beats ALWAYS NO, a scenario so rare that its classical probability is zero.² One might object that little intuitiveness is lost if (3) is replaced with:

- (6) One cannot use past information about the rolls of a memoryless die to get a *significantly* better guessing strategy than the best strategy that does not use past information.

A strategy that does better in an extremely rare case is not significantly better.

But now note that one can do even better than ALMOST ALWAYS NO. The thought behind ALMOST ALWAYS NO was that there is a set S of sequences of rolls such that

- (7) you can tell based on infinite past information that the infinite sequence of rolls falls in C ,

² *Suppose the rolls are numbered. Let U_n be the event that all rolls prior to roll number n are sixes. Then $P(U_n) = (1/6)(1/6)(1/6) \dots = 0$. Let U be the event that there are only finitely many non-sixes. Then $U = \bigcup_n U_n$. By countable additivity, if $P(U_n) = 0$ for all n , then $P(U) = 0$. And, intuitively, even if we do not have countable additivity, we would not think U would have a probability greater than an infinitesimal.

and

- (8) once you know that a sequence of rolls is in C , you can ensure there are only finitely many mistakes in your guessing.

The set that ALMOST ALWAYS NO leveraged was C_6 , the set of roll sequences that are almost all sixes. We can improve on ALMOST ALWAYS NO by adding to the strategy special cases for other sets satisfying (7) and (8).

For instance, we can add D , the set of all sequences of rolls such that almost all the odd-numbered ones are sixes and almost all the even-numbered ones are non-sixes. We can tell on the basis of past data whether the sequence of rolls that one is observing is a member of D , and then one can guess in accordance with D : guessing six for odd-numbered rolls and non-six for the even-numbered ones. Adding this rule to ALMOST ALWAYS NO means that we will do better than ALWAYS NO both for sequences in C_6 and for sequences in D . Again, the improvement is insignificant, however, since the classical probability that a given sequence of rolls is a member of D can be shown to be zero.

The strategy can be extended. We can come up with a set V of disjoint sets of sequences of rolls such that each member C of V satisfies (7) and there is a strategy T_C that when followed guarantees there will only be finitely many errors if the sequence of rolls is in C . Then we have the following strategy:

- (9) EVEN BETTER: If almost all rolls are non-sixes or the sequence of rolls doesn't fit in any member of V , guess "Not a six"; otherwise, let C be the member of V that contains the sequence and guess according to T_C .

ALMOST ALWAYS NO had $V = \{C_6\}$, where T_{C_6} was the strategy of always guessing "Six", and our slightly improved version had $V = \{C_6, D\}$, where T_D was the alternating guess strategy.

Yuvay Gabay and Michael O'Connor (see Hardin and Taylor 2008) used the Axiom of Choice basically³ to prove that one can make a set V large enough that *every* sequence of rolls falls into some member of V and there is a strategy T_C for every member C of V guaranteeing at most finitely many errors. One can then tweak the strategy to guarantee that if almost all rolls are non-sixes, then one always guesses "Not a six". (See Section 2.4.) If we do that, then EVEN BETTER will be equivalent to ALWAYS NO if almost all rolls are non-sixes, but for all other sequences of rolls it will be better, as EVEN BETTER will make only finitely many mistakes, while ALWAYS NO will make infinitely many. This certainly counts as a *significant* improvement in the sense of (6). For, except in the extremely unlikely case that all the rolls are non-sixes (the classical probability of that case is zero), we improve from infinitely many electric shocks to finitely many.

³ Their setting involved guessing hat colors. See also Thorp (1967) for an earlier but somewhat different solution.

So, given causal infinitism, we can leverage past data to do significantly better than ALWAYS NO. That is absurd, so we should reject causal infinitism.

There is, however, one loose end. The strategy EVEN BETTER depends on the Axiom of Choice, specifically in the choice of strategy S_C for each member C of V . The details of this dependence will be given shortly for the technically minded reader, but there is a real worry about whether a strategy “generated” by the Axiom of Choice could be causally implemented. In Chapter 6, I will argue, however, that such strategies can be causally implemented if causal infinitism is true (I will also argue for the truth of a version of the Axiom of Choice sufficiently strong for the mathematics of the present argument). That argument will close the loose end in the present *reductio ad absurdum* against causal infinitism.

2.4 *Construction of strategy guaranteeing at most finitely many errors

Let Ω be a collection of backwards-infinite sequences $\dots, a_{-2}, a_{-1}, a_0$ of numbers from $\{1, 2, 3, 4, 5, 6\}$. Let \sim be the equivalence relation (transitive, symmetric, and reflexive relation) on Ω where $x \sim y$ if and only if x and y differ in at most finitely many places. Let V be the set of all equivalence classes $[x] = \{y : y \sim x\}$ for $x \in \Omega$. By the Axiom of Choice, there is a choice function for S , i.e., a function f on V such that $f(C)$ is a member of C for every $C \in V$. We can also impose one further condition on f : if $C \in V$ has the property that no member of C contains infinitely sixes, then $f(C)$ has no sixes at all. To do this, for such a C , simply replace every occurrence of 6 in $f(C)$ with an occurrence of 1. There will only be finitely many replacements, so we will stay within the equivalence class C .

Now, for any $C \in V$, the strategy T_C is very simple: if $f(C)$ is 6 in the n th place, guess “Six” before the n th toss; otherwise, guess “Not a six”. Every member of C differs in at most finitely many places from $f(C)$, so this strategy will result in only finitely many mistakes. Furthermore, the further condition imposed on f ensures that sequences consisting of almost only sixes lead to guessing “Not a six”.

2.5 A multipersonal synchronic version

2.5.1 AN ANGELIC ANNOUNCEMENT

Infinitely many perfectly rational people each roll an independent and fair die. Things are arranged so they can’t see how the rolls come out, and yet each is asked to guess if her roll came out six. If a person guesses correctly, she gets a treat; otherwise, she gets an electric shock.

It’s obvious what each person should do: She should guess “Not a six”. But given causal infinitism, an angel could know how all the rolls came out. And it’s possible that almost all rolls—i.e., all but at most finitely many—come out sixes. In that case, we suppose the angel announces to all that almost all rolls came out six. We also suppose that all the persons are certain that the dice are independent and fair, and that the angel only tells the truth.

If you are one of the die rollers, then before the announcement you were planning to guess “Not a six”. Should you change your guess?

There are convincing arguments for both options. In favor of changing to “Six” is the simple thought that only finitely many rolled anything other than six, and so it seems overwhelmingly likely that you rolled a six. How likely is it, after all, that you are one of the finitely many exceptions? No more likely than anyone else, and surely it’s unlikely you had won a countably infinite lottery with only finitely many winners.

Furthermore, if everyone switches to “Six”, then only finitely many electric shocks will be administered, while if everyone sticks with “Not a six”, almost everyone will be shocked. It seems, thus, that the right answer had better be that everyone should switch to “Six”.

Granted, there are famous cases where everyone’s doing what is in their own interest results in a worse overall outcome. The most famous is the Prisoner’s Dilemma. Two prisoners are captured, each unable to communicate with the other. If both keep silent, only minor charges can be pinned on them, and they will each get a year in jail. On the other hand, if one keeps silent and the other informs (“defects”), the silent one will get ten years in jail and the informer will be set free. But if both inform, then they will each get nine years in jail, since major charges will be proved against each, but their informing will make the verdict more favorable. Now, the crucial point about a Prisoner’s Dilemma is that no matter what the other prisoner does, if the only thing that is at stake is avoiding jail—thus, neither morality is at stake nor are there going to be any future dilemmas—it is better to defect. For if the other prisoner defects, one’s own defection will shorten one’s sentence by a year, while if the other prisoner is silent, one’s own defection will get one off completely. But if each does this apparently most rational thing, both end up in jail for nine years, which is a much worse outcome than if both remain silent. So if each does what is most rational, the outcome is worse overall for each than if each is silent.

Now, it is controversial whether or not it is rational to defect. But if the rationality of defection is granted, then we do indeed have the result that each person’s doing what is in her own self-interest can produce overall harm. This conclusion is pretty intuitive. We get a similar structure in the Tragedy of the Commons. If everybody grazes their sheep on the commons, the commons will be overgrazed and there will be but little benefit to each person. But each person still gains by grazing her own sheep on the commons: grazing on a depleted commons is better than not grazing on the commons at all, or so the story goes.

However, notice that in these standard cases, each person’s action affects other people. Informing on a fellow prisoner changes the other’s sentence, and grazing one’s sheep decreases the grass available for other people’s sheep.

Another family of cases where everyone’s acting rationally in their own self-interest is harmful is when people are ignorant of relevant factors. If everyone knows that cyanide prevents cancer, but does not know that it prevents it by killing the patient, then everyone’s acting rationally in their own self-interest results in the extinction of humanity.

However, it is very plausible that when people’s actions affect only their own well-being, each person’s rational evaluation of what is good for her should align with

her evaluation of what would be good for everyone to do in light of everyone's well-being. So given that each person can see that on balance things would go better if everyone guesses "Six", this gives us good reason to think that that is the individually rational thing to do.

On the other hand, there is a strong argument that one should stick with the guess "Not a six". For one should not change one's rational guess upon receipt of epistemically irrelevant information. Information about other people's rolls is irrelevant when the dice are independent. Therefore one should not change one's guess upon learning:

(10) Among people other than yourself, almost everyone rolled six.

(Of course, it is crucial that you are *certain* that the dice are fair and independent. Without that certainty, (10) would convince you that the dice are unfair.) But (10) is logically equivalent to:

(11) Almost everyone rolled six.

For your roll cannot affect whether infinitely many people other than you rolled a non-six. Since (10) shouldn't change what it is rational to guess, neither should the equivalent (11).

The above arguments show that in the situation described, one should switch to guessing "Six" and that one shouldn't switch to guessing "Six". Assuming there are no possible real rational dilemmas—cases where one rationally should do something and one should rationally refrain from it—this is impossible. So something is impossible in our scenario. Causal finitism provides an elegant explanation of what that is. The angel's announcement, in order to be absolutely trustworthy, would need to be causally responsive to the infinitely many die rolls. But that is impossible given causal finitism.

Or, to put it in the usual form of our argument, if causal infinitism is true, the above die-rolling scenario is possible. But if it's possible, it's possible that it be rationally required that one guess "Six" and it's rationally required that one not do so. But that is impossible, so causal infinitism is false, and hence causal finitism is true.

2.5.2 AN OBJECTION AND A TWEAK

One of the arguments for switching to "Six" after the angel's announcement was based on the idea that if the number of non-sixers is finite, you will judge it very unlikely that you are a member of the set of non-sixers.

This particular argument, however, may need some refinement. Very plausibly, if a subset N of rollers is produced in a way that isn't biased in favor of any rollers, say by letting it be the subset of all who didn't roll six, and the cardinality of N is announced, then when that cardinality is infinitely smaller than the cardinality of the whole set of rollers, your credence that you are a member of N should be very small. It follows that if a particular finite cardinality is announced for the set of non-sixers, you should switch your guess to "Six". But that isn't the story I gave: what was announced was

that N is finite, and “finite” doesn’t denote a particular cardinality (there are infinitely many finite cardinalities, namely $0, 1, 2, \dots$). Nonetheless, I think it is highly intuitive that the credence that you are in N should still be zero or at least nearly zero.⁴

Still, for the sake of readers impressed by this concern, one can modify the case as follows. We first stipulate that the cardinality of the number of rollers is not merely infinite, but is an uncountable infinity, say the cardinality c of the continuum, that is infinitely many times greater than the countable infinity \aleph_0 of natural numbers.⁵ Then, let’s suppose that instead of the number of rollers who rolled non-six being finite, it is countably infinite, and the angel announces that there are \aleph_0 non-sixes. Then even if one has qualms about the difference between an announcement that the cardinality is smaller and an announcement of particular smaller cardinality, one should conclude here that the credence that one rolled non-sixes will become nearly zero.

The other arguments all adapt neatly to this modified version of the story. In particular, the crucial observation that (12) and (13) are logically equivalent can be replaced by the observation that the following two claims are logically equivalent:

- (12) Among people other than you, all but \aleph_0 rolled six.
- (13) All but \aleph_0 people rolled six.

For, just as a single roll doesn’t affect whether there are finitely or infinitely many sixes, a single roll doesn’t affect whether there are \aleph_0 or not.

2.5.3 *MAKING THE PARADOX ROBUST

The interpersonal version of the guessing paradox so far only works in the very unlikely—but possible!—case where almost all people roll a six. But just as we used the Axiom of Choice to make the one-person sequential version work no matter how the dice come out, one can do the same here. I leave it to the reader to decide whether this makes the paradox more compelling: I am not sure anything is gained by the extra complication.

⁴ *Ian Slorach points out to me that the issues involved with an announcement of “finite” are connected with conglomerability (see Chapter 4, Section 2.5.3). For any particular finite cardinality, if that cardinality were announced for N , your credence plausibly should be nearly zero. But it requires conglomerability to conclude from this that if it’s merely announced that the cardinality is finite then your credence plausibly should be nearly zero, and in these kinds of infinitary situations we cannot expect conglomerability. I agree that conglomerability would be needed if the inference that the credence that one is in N is nearly one were made from the fact that given any particular finite cardinality the credence that one is in N would be nearly zero. But I think there is a *direct* intuition that if N is a finite subset of an infinite set and chosen in an unbiased way, you are very unlikely to be in N .

⁵ We could say that an infinite set B is infinitely many times larger than an infinite set A if B can be partitioned into infinitely many subsets each of which is larger than some fixed cardinality κ which is greater than that of A . And the real numbers, whose cardinality is the continuum, can be partitioned into infinitely many subsets of continuum cardinality, say the intervals $\dots, [-1, 0), [0, 1), [1, 2), [2, 3), \dots$, and the cardinality of the continuum is greater than the cardinality of the natural numbers.

The method is pretty much the same as before. One defines the equivalence relation on possible outcomes (where an outcome can be thought of as a function from persons to the set $\{1, 2, 3, 4, 5, 6\}$) by saying that two outcomes are equivalent provided that they differ in at most finitely many places. A choice function that chooses a particular outcome $f(A)$ in each class A of equivalent outcomes is fixed. Then as soon as the dice are rolled, the angel checks which class A of equivalent outcomes the actual outcome is in, and then announces $f(A)$. It is now guaranteed that almost all rolls match the outcome $f(A)$. Again, we have good arguments that each agent should guess in accordance with $f(A)$. Yet learning that the outcome of all the rolls is equivalent to $f(A)$ tells you nothing about your own roll, since your own roll doesn't affect whether the overall outcome matches $f(A)$ for almost all people.

This strategy guarantees that almost all guesses will be right, and yet each guess still depends only on the other rolls. There is, however, a disanalogy between the synchronic and sequential cases. In the sequential case, I claimed that it was surprising that given information on an infinite amount of *past* data it was possible to even guarantee a single correct answer (see the Theorem in Section 2.1). In the unordered synchronic case, that one can guarantee some correct answers when each is supplied information about others' rolls is not so surprising.

For instance, with two people rolling dice, and each finding out the other's roll, one can guarantee that at least one guesses correctly: for instance, suppose Alice guesses in accordance with Bob's roll ("Six" if six and "Not a six" otherwise) while Bob guesses in opposition to Alice's roll ("Six" if not a six and "Not a six" if six). Then when they roll the same, Alice gets it right; when they roll differently, Bob gets it; and so always exactly one gets it right. But notice that the expected overall outcomes from this strategy is one shock and one treat, whereas both guessing "Not a six" has the expectation of 0.33 shocks and 1.67 treats, so that the strategy is overall inferior to sticking with "Not a six".

2.6 A parody?

2.6.1 THE STORY

But there appears to be a parody of our reasoning. As before, we suppose that countably infinitely many people roll a die but don't see the result, and they must guess whether they rolled a six. Again, they get a shock if they get it wrong, and a treat if they get it right. Moreover, every person is a complete stranger to every other. Clearly, if that's the whole story, everyone should guess "Not a six".

Now, it is nearly certain—i.e., has probability one or one minus an infinitesimal—that infinitely many of the rolls will be sixes and infinitely many will be non-sixes. Suppose that this happens. Prior to anyone's having a chance to make any guesses, an angel then divides the people into groups of three, where each group of three contains exactly two sixers—people who rolled a six—and one non-sixer,⁶ and transports

⁶ Suppose the sixers are s_1, s_2, \dots and the non-sixers are n_1, n_2, \dots ; then the angel can form groups corresponding to the sets $\{s_1, s_2, n_1\}, \{s_3, s_4, n_2\}, \{s_5, s_6, n_3\}, \dots$

each such trio into a separate room. Everyone is made certain that the angel has done this.

If you're in one of the rooms, you know that of the three people in the room, exactly two rolled sixes. How should you guess? There is a fine argument that you learned nothing. You already knew you'd be transported into a room with two complete strangers, and that is exactly what happened. Having learned nothing, you should stick to "Not a six". On the other hand, there is something absurd about all three people in the room guessing "Not a six" when it is certain that two of them got sixes. Within each room, it is surely better that the three people each guess "Six". So once again we have a good argument that each person should guess "Six" and that each should guess "Not a six".

However, there is a crucial difference between the rearrangement story and our previous story about guessing almost all rolls. In our previous story, embracing causal finitism killed the possibility of the paradox. However, in the present story it is not clear that any infinite causal history is involved.

For instance, consider a world where exactly one new human being is conceived (which I will assume for simplicity is the beginning of her existence) at the beginning of each year, starting with year 1, and no one ever dies. At age 30, each human rolls a die. Once she has rolled a die, the die is taken away by the angel, without the human seeing the result, and the person is put into a dreamless sleep. The angel now apportions people to rooms sequentially, by keeping track of the sequential die throws, and filling a room as soon as two sixers and one non-sixer is available. When a room is filled, the people in it are woken up, with no sign of how much time has gone by,⁷ and then they get their opportunity to guess whether they rolled a six.

Completing this process will take an infinite amount of time. But nearly certainly (the exception is the unlikely case that almost everyone rolls a six or almost everyone rolls a non-six) it will only take a finite amount of time for any one person to be placed in a room. Finitism *simpliciter* plus eternalism will rule out the scenario, by ruling out the possibility of infinitely many future events, but causal finitism is compatible with the scenario. We thus have a parody: a story apparently paradoxical for the same sorts of reasons as our previous die-guessing stories, but without any dependence on causal infinitism.

2.6.2 EVALUATING THE PARODY

To see how powerful the parody is, we need to see how closely analogical the arguments for staying with "Not a six" and for switching to "Six" in the rearrangement case are to the almost-all-six case.

In Section 2.5.1, we considered two arguments that after finding out that almost everyone rolled six you should switch to guessing "Six". The first was an intuitive one: You could equally well be any one of the people, and it is very unlikely that you would

⁷ Knowing how much time went by would be a clue to whether one rolled a six or not.

be one of the few who didn't roll six. This argument is closely analogous to the thought that if two of the three people in your room are sixers, you have a $2/3$ chance of being a sixer.

But as far as it goes, the latter thought is mistaken. Take a variant where there are a million people, and where the angel has antecedently promised you, and only you, to put you in a room where there are two sixers and one non-sixer. Then while it's still true that of the three people in the room exactly two are sixers, you shouldn't conclude you have a $2/3$ probability of being a sixer. Your chance of being a sixer admittedly goes up insignificantly since by being put in such a room, you now do learn for certain that there are at least two sixers and at least one non-sixer, whereas previously you only thought that this was extremely likely to be true. But apart from that piece of information, nothing relevant has been learned, certainly nothing to boost your probability to $2/3$. So the mere fact that in your room there are two sixers and one non-sixer doesn't imply a $2/3$ probability of being a sixer.

But perhaps in the infinite case we can say that the lone non-sixer is equally likely to be any one of the three people in the room, as the case is fully symmetric between the three people. In the finite case there is no such symmetry. Your presence in the room is guaranteed by the angel, while the other two people are picked randomly. Similarly, the case where almost everyone rolls six was symmetric between the participants. In both cases, then, intuitive considerations of symmetry suggest that you should switch to "Six".

But the symmetry considerations have some weaknesses in the original parody case. I suggested that the parody case could be run without causal infinitism by having the angel assign people to rooms by an ordering determined by conception order and die roll results. Now if you knew the conception dates of all three people, that would give you significant information as to who rolled six and who did not. Indeed, the oldest person in the room is most likely to be the one to have rolled a non-six. For the angel goes through the list of sixers—starting with the oldest—twice as fast as through the list of non-sixers, since the angel has to put two sixers into a room for each non-sixer, and moreover the sixers are on average more widely spaced in age. So if you knew all the relevant necessary truths, the situation would no longer be symmetric to you.

In order to maintain symmetry, you must have no data at all distinguishing you from other people with respect to conception order. But that means that your position in order of conception is, at least from your point of view, the value of a countably infinite fair lottery. Now we have two plausible stories here. First, we have an essentiality of time of origination thesis that ensures that you couldn't have had a different position in the conception order than you did. Second, you could have equally had any of the infinitely many positions in the conception order.

Given the essentiality of time of origination thesis, it is a necessary truth that you came into existence at the time you did, and that the other people in your room came into existence at the times that they did. This is disanalogous to the original case

where there is a symmetry between the die rollers, even after fixing their identities and ages. Granted, there is still an *epistemic* symmetry in the parody case. But when an epistemic symmetry is due to ignorance of necessary truths, there is good reason to be worried about the correctness of one's probability assignment, since in such a case one is assigning non-unit epistemic probability to relevant necessary truths, which makes one's probabilities be relevantly inconsistent, and it is neither surprising nor paradoxical if this leads to strange consequences.

On the other hand, if you could equally have had any of the infinitely many positions in the conception order, then the process which led to the conception order implemented a countably infinite fair lottery—maybe the winner is the first one conceived—and we saw in Chapter 4 that such lotteries are impossible.

One might try to order the die rollers in some other unbiased way than by conception order. But to do that seems to require either ignorance of necessary truths or countably infinite fair lotteries.⁸

In our interpersonal guessing case, the second argument for switching was based on a universalization principle that, roughly, if an action that doesn't affect others is self-interestedly rational, its generalization should produce the best overall results. This argument does not generalize to the parody case. For, nearly certainly, if everyone switches to "Six", infinitely many people will get a shock and infinitely many will get a treat. And it will be exactly the same overall if everyone sticks with "Not a six": infinitely many shocks and infinitely many treats. Granted, the best consequence for the agents *in your room* will obtain from the three occupants all switching to "Six". But the universalization principle is less plausible when restricted in an arbitrary way to the denizens of the room.

So of the two arguments for switching in the parody case, one is damaged by essentiality of time of origin, and the other is based on a less plausible universalization principle.

What about the argument for staying with "Not a six"? Our original argument was that the proposition that almost everyone rolls six is logically equivalent to the proposition that almost everyone other than you rolls six, and since the latter is evidentially irrelevant (given certainty about the fairness of the setup), so must the former be. This argument does not apply in the parody case.

⁸ *The Sleeping Beauty problem (Elga 2000) is similar to our parody on an essentiality of time of origins interpretation, in that in Sleeping Beauty it is essential that the agent be ignorant of which day of the week it is when she wakes up, and it is a necessary truth that today is Tuesday if, in fact, today is Tuesday. Ignorance of such relevant necessary truths can be expected to lead to strange results. However, in Sleeping Beauty ignorance of necessary truths is not essential to the story. While the standard story involves a selection between wakeups on Monday and Tuesday versus a wakeup only on Monday, so that knowledge of the necessary truth that it is now Tuesday tells one how the selection went, one can tell a modified story where the selection is between wakeups on Monday and Tuesday versus a wakeup on a day randomly chosen from among Monday and Tuesday, in which case knowledge of the necessary truth that today is, say, Tuesday will not help.

However, there is an additional argument that can be given for sticking to “Not a six” in the parody case. First, there is the plausible thought that you learn nothing relevant by being put in the room with the two strangers. After all, you already knew that you *could* be put in such a room (assuming that there are infinitely many sixers and infinitely many non-sixers), and the identities of the people in the room give you no relevant information.

This argument is deficient, however. For given the essentiality of time of origin, and given that the angel apportions people to their rooms in order of coming into existence, if you knew all the relevant necessary truths, you would reassess your probabilities. But we need to be cautious about our probabilistic judgments where we are ignorant of relevant necessary truths, as then our probabilities are bound to be inconsistent.

All in all, the arguments for both sides of the dilemma in the parody story are weaker than in the almost-all-sixes story, and disanalogous to it despite surface similarities.

I do not know what the right resolution to the parody problem is. Here is one option, however. The story involves either a countably infinite fair lottery implied by a random ordering, and those are impossible by the considerations of Chapter 4, or ignorance of necessary truths such as time of origin. It was crucial to the guessing paradoxes that one can’t be rationally required to do something and to refrain from doing it. But perhaps this principle only applies to perfectly rational beings, or at least beings that are perfectly rational in respect of the relevant features of the situation. Maybe just as some (e.g., Aquinas; see Dougherty 2011, Chapter 4) think you can get yourself into a real moral dilemma through moral imperfection—say, by making incompatible promises to different people—you can get yourself into a real rational dilemma through rational imperfection. And perhaps ignorance of relevant necessary truths—such as the year that one originated in—makes it impossible to be perfectly rational about an area of thought.

3. Satan’s Apple

3.1 *The story*

Arntzenius and Hawthorne (2004, Section 3) give this paradoxical situation:

Satan has cut a delicious apple into infinitely many pieces, labeled by the natural numbers. Eve may take whichever pieces she chooses. If she takes merely finitely many of the pieces, then she suffers no penalty. But if she takes infinitely many of the pieces, then she is expelled from the Garden for her greed. Either way, she gets to eat whatever pieces she has taken.

(The pieces, of course, must get thinner and thinner.) There are two versions of this: a synchronic version where Eve simultaneously decides which pieces to take and a diachronic version where Eve is presented, one by one, with a binary choice whether to take a given piece.

3.2 *Synchronic version*

First take the synchronic version. Here, Eve needs to indicate which subset of pieces she wishes to eat. The reason paradox appears here is that it seems that for each slice she is better off including that slice in her indicated subset, no matter what she does about the other slices (better be expelled from the Garden with an extra slice than without it). This suggests a dominance argument: for each slice, she should choose to include that one, because no matter what other choices she makes, it's better to include it.

As Arntzenius and Hawthorne (2004) note, this kind of dominance reasoning works in the finite regime. If for each of a finite number of slices you're better off including that in your selection, no matter what you choose about the other slices, then you should go for all the slices. The reason is this. In a finite case, there will be at least one optimal option (there may be more than one, in case of a tie). If that optimal option failed to include a slice, then it wouldn't be optimal since there would be a better option, one that includes that slice.

In general, the dominance argument establishes that *if* there is an optimal solution, it includes all the slices. But in Satan's Apple the antecedent of the conditional is false. For each profile that is missing a slice of the apple, including that slice would be an improvement. And the profile that includes all the slices is beaten by the profile that includes none of them: better stay in paradise than eat the whole apple. Hence no profile is optimal.

There is another reason why the dominance argument is problematic. It conflates the synchronic and diachronic cases by making it seem like Eve is making an infinite number of choices, one per slice. But she is instead making a single choice between an infinity of options.

Granted, it would be nice to have an optimal solution to the problem. There is something disquieting about the fact that for each solution there is a better. It would be nice if causal finitism ruled out such scenarios. Alas, I do not know how credibly this can be done.

One possibility would be to argue that agents deliberate between multiple reasons, and these reasons compete causally in the agent, with one of the reasons ending up being the winner—for instance, this is Kane's (1999) picture of human freedom. But then it seems that Eve's choice would have to be the output of a causal process with infinitely many reasons as antecedents, there being a reason corresponding to each slice.

I think this response to the synchronic problem fails on two counts. First, multiple reasons could correspond to a single state of the agent. It is plausible to attribute to someone who believes that the moon is round and gray the belief that the moon is round as well as the belief that it is gray. But we need not think the mind stores all three beliefs separately. It could be that there is a single mental state which grounds the correctness of all three belief attributions. Likewise, an agent can have many reasons and a single mental state that grounds the attribution of these

reasons. In that case, there may be only one causal factor involved, namely that mental state.

Second, even if each reason corresponds to a different causal factor, it seems that only one reason—the victorious one, the one that produced the action—is actually in the causal history of the action. The competing reasons were counter-explanatory: rather than contributing to the action, they hindered it.

There may be some other way for causal finitism to resolve the synchronic problem. But we can also just agree with Arntzenius and Hawthorne (2004) that there is no paradox. This is just a case where there is no optimal option (this is a way of solving rather than killing the paradox, in the terminology of Chapter 1, Section 1). And anyway we shouldn't expect causal finitism to resolve *all* problems that involve infinity.

3.3 *Diachronic version*

The diachronic version is more paradoxical. Arntzenius and Hawthorne (2004) divide the case depending on whether Eve can bind herself ahead of time to a particular pattern of future decisions. If she can bind herself, the problem is equivalent to the synchronic one: she just needs to choose a profile of slices and bind herself to it. We still have the difficulty that there is no best option, of course, but perhaps that's not a real paradox.

But the really difficult case is where Eve cannot bind herself ahead of time. Here, Arntzenius and Hawthorne have to admit that it will be rational for Eve to take each slice, even if this damns her.

The diachronic version, however, can be neatly ruled out by causal finitism. For Eve's expulsion from the Garden would have to be caused by infinitely many causal factors, presumably arranged in a supertask. So even if causal finitism does not help with the synchronic problem, it kills the diachronic paradox.

We can also imagine a multipersonal synchronic version that has the same features. Infinitely many people are in paradise and have a choice whether to eat an apple. If infinitely many eat it, they are all expelled. If finitely many eat, they all stay. What should they do? Again, each is better off eating an apple, no matter what the others do. But if they all eat an apple, a terrible thing happens. Again, causal finitism neatly kills the paradox by making the story impossible.

3.4 *Objection: Scores, desires, and promises*

It was crucial to the causal finitist resolution of the diachronic paradox that the payoff causally depended on an infinite number of events.

But payoffs don't need to depend *causally* on antecedents. Suppose we play a guessing game where I toss a coin you don't get to see and then pay you a dollar if you guess correctly. In that case, the payoff causally depends on my checking whether the outcome matches the victory condition. But then we tweak the game. I toss a coin and you guess how it came out, and you get victory, and victory alone, if you guessed right. You don't get any money, and you don't even find out that you won. All you have

is victory pure and simple. In that case, the payoff doesn't depend on anyone's being able to check how the coin landed.

If victory alone is worth pursuing, then so is a good score in a game, even absent knowledge of that score. One can pursue running a mile quickly even when one doesn't expect to find out how long one took to run. But this allows us to modify Satan's Apple. For instance, we can suppose a game where eating slice n wins you $1/2^n$ points but eating an infinite number of slices loses you 100 points.

There is, however, an intuitive difficulty here. Victory is not much fun if no one finds out you won, and by the same token a high score isn't worth much if no one finds out you got it. And finding out the score would depend on the infinite number of decisions.

In fact, there is good reason to doubt that victory and score generate *any* real benefit to an agent, in and of themselves. Consider this solitary game: I guess "even" or "odd". If at the time of my guess the number of mosquitoes in the world matches my guess, I get two points. Otherwise, I lose a point. I can play this over and over, winning roughly every second time. If it's good for me to win at a game, I continue to rack up benefits. So for reasons of self-interest, I should play this game all the time. I could even set myself up as playing it by default: I stipulatively announce that my breaths alternate between "even" and "odd" guesses. I will be racking up benefits every day, every night. But that's silly.

But perhaps little benefit accrues to the agent in this silly game because it is unbalanced—little effort is needed to have a positive expected score—and maybe there are diminishing returns on replaying a game, so that even an infinite number of wins has but little benefit. However, it is also plausible that there is *no* benefit in the game, that we have no reason to play the game.

Here, however, is a way to resolve the conflict between the intuition that one can reasonably play for score alone even when one knows one won't find out the score and the intuition that there is no point to the mosquito guessing game. Perhaps it is essential to a game that it be the sort of practice where normally we could find out the score (or at least a good approximation to it). There can be individual cases where one benefits from or is harmed by an unknown score—perhaps on your morning run around the block unbeknownst to you you beat the world record for a 400-meter run, and if so, bully for you. But these cases are plausibly parasitic on games of the same, or a very similar, sort where the scores are knowable, unlike the gamified Satan's Apple.

Moreover, even if the score is worth having for its own sake, the mapping between scores and utilities—i.e., measures of well-being—is unlikely to be simple, and is not going to be a subject for mere stipulation. We could make a variant of soccer where we stipulate that a goal scored in the last minute of a game counts for a thousand points. But a team that won by a thousand to zero wouldn't be a thousand times better off than a team that won by one to zero, no matter how much we tried to stipulate this. And while an overwhelming victory is worth more than a slight victory against the same opponents, it is not clear that winning a basketball game by 83 to 15 is any more

valuable than winning it by 82 to 15. Nor is it clear that running a mile in 4.000002 seconds is worse for one than running it in 4.000001 seconds, especially if one doesn't know the time.

In light of this, there is no guarantee that it is actually possible to set up a game with utilities that yield a successful gamification of Satan's Apple—a scenario where choosing infinitely many slices gives one the worst score, but adding one slice is always beneficial. Any utility we get from the mere score of a game is surely finite, and it is unclear that the scores-to-utilities conversion can be so fine-grained as to be able to produce arbitrarily small increments of utility.

Finally, even if there really are benefits from silly games with unknown scores like the gamified version of Satan's Apple, perhaps the existence of paradoxes in *pure games*, games where score alone is the value, is not that surprising. Such pure games are not something that contributes a significant part of human well-being, and just as one might think that our intuitions about probabilities are insensitive to infinitesimal differences (cf. Chapter 4, Section 2.6), so too one might think that our intuitions about rationality may not apply to pure games. In any case, where there are "concrete" benefits—such as healthy pleasure and the avoidance of unhealthy pains, as in the original version with a tasty apple and a pleasant garden—the paradoxes seem more significant.

A variant on generating utilities non-causally is doing so by means of desires. It seems worthwhile to have one's desire satisfied, even if one never finds out about the satisfaction. A painter may desire her painting to hang in the National Gallery. She benefits from its hanging there even if she never finds out that it is there.

Just as we considered using game scores to power Satan's Apple, we could try to make the utilities turn on desires. We could imagine an agent who has an overwhelmingly strong desire not to eat an infinite number of apple slices, but who nonetheless desires each particular slice, in such a way that it's always better to have one more slice, but having infinitely many is worse than any scenario with finitely many.

Similar responses can be made to the desire version as to the score version. First, it is unclear whether satisfaction of a desire by itself, in the absence of a desire-independent good being desired and in the absence of awareness of the satisfaction of the desire, is valuable. Second, it may be that the value of unconscious satisfaction is parasitic on the value of conscious satisfaction, so that only the sorts of desires where one can find out about the satisfaction are valuable. Third, it is not clear that mere satisfaction of desire can be sufficiently fine-grained to run the paradox. And, finally, it is not clear that mere desire satisfaction counts for enough to make the paradox seriously problematic—again, perhaps the intuitions about rationality are not sensitive to slight differences in utility.

3.5 *Evaluation*

We can resolve the original diachronic Satan's Apple by means of causal finitism. The synchronic version can be resolved differently, and non-causal versions based

on games, desires, and promises do not appear compelling. I do not know if *every* variant of Satan's Apple can be resolved by a combination of causal finitism with other tools. But the fact that some versions can be resolved with causal finitism, and other known versions are resolved with other tools that do not resolve the version that is resolved with causal finitism, gives us some reason to accept causal finitism, though not as much as if we could resolve all versions uniformly with causal finitism. Still it is somewhat impressive that causal finitism helps with the most compelling version of the paradox.

4. Beam's Paradox

Beam (2007) gives a more complicated betting paradox where you should accept each of an infinite collection of deals, but where accepting all the deals guarantees a net loss. Again, this paradox has the feature that it can be implemented in multiple ways, some of which can be resolved without causal finitism, but the most troubling version seems best resolved by means of causal finitism.

4.1 **The mathematical formulation*

Start by observing that there is a permutation π of the positive integers such that:

$$\sum_{n=1}^{\infty} \frac{(-1)^{\pi(n)}}{\pi(n)} = -100.$$

This follows from the fact that $\sum_{n=1}^{\infty} (-1)^n/n$ is only conditionally convergent, and a conditionally convergent series can be rearranged to have any value one wishes.

Let X be a uniformly randomly chosen point in the interval $(0, 1)$. Define $a_n = (-1)^{\pi(n)+1}$.

Assume dollars measure utility. For each n , in exchange for being paid $(1/2)^{\pi(n)}$ dollars, a rational agent will play this subgame:

- (i) If $X < 1/\pi(n)$, get $(1 - 1/\pi(n))a_n$;
- (ii) Otherwise, get $-(1/\pi(n))a_n$.

The expected value of this game is

$$\begin{aligned} \left(1 - \frac{1}{\pi(n)}\right) a_n P\left(X < \frac{1}{\pi(n)}\right) - \frac{1}{\pi(n)} a_n \left(1 - P\left(X < \frac{1}{\pi(n)}\right)\right) \\ = \left(1 - \frac{1}{\pi(n)}\right) a_n \frac{1}{\pi(n)} - \frac{1}{\pi(n)} a_n \left(1 - \frac{1}{\pi(n)}\right) = 0. \end{aligned}$$

Thus one rationally should play the subgame no matter how little one is paid to do so, including when one is paid $(1/2)^{\pi(n)}$ dollars.

Suppose you accept the deal for each n . Let us compute your net outcome for a given value of X . The unconditionally convergent series $\sum_{n=1}^{\infty} 1/2^{\pi(n)} = \sum_{n=1}^{\infty} 1/2^n = 1$

represents how much you will be paid just to play the subgames. Let $A_n(X) = \{n \geq 1 : \pi(n) < 1/X\}$ and $B_n(X) = \{n \geq 1 : \pi(n) \geq 1/X\}$. Then for $n \in A_n(X)$, rule (i) will apply and for $n \in B_n(X)$, rule (ii) will apply. Moreover, the set $A_n(X)$ is finite no matter what $X \in (0, 1)$ is, and even in a conditionally convergent series one can rearrange *finitely* many terms. For brevity omitting the dependence on X , your payoff from the subgames will be:

$$\begin{aligned}
 \sum_{n \in A_n} \left(1 - \frac{1}{\pi(n)}\right) a_n + \sum_{n \in B_n} -\frac{a_n}{\pi(n)} &= \sum_{n \in A_n} (-1)^{\pi(n)+1} \left(1 - \frac{1}{\pi(n)}\right) + \sum_{n \in B_n} \frac{(-1)^{\pi(n)}}{\pi(n)} \\
 &= \sum_{n \in A_n} (-1)^{\pi(n)+1} + \sum_{n \in A_n} \frac{(-1)^{\pi(n)}}{\pi(n)} + \sum_{n \in B_n} \frac{(-1)^{\pi(n)}}{\pi(n)} \\
 &= \sum_{n \in A_n} (-1)^{\pi(n)+1} + \sum_{n=1}^{\infty} \frac{(-1)^{\pi(n)}}{\pi(n)} \\
 &= \sum_{n \in A_n} (-1)^{\pi(n)+1} - 100 \\
 &= \sum_{1 \leq m < 1/X} (-1)^{m+1} - 100
 \end{aligned}$$

dollars. But $\sum_{1 \leq m < 1/X} (-1)^{m+1}$ equals either 1 or 0, depending on whether there is an odd or even number of terms in that sum. In either case, you will lose at least \$99 in playing, and since you will have been paid \$1 to play, your net loss will be at least \$98.

Hence, if you do what is rationally required, you are guaranteed to lose at least \$98. The crucial ingredient in the paradox is that in such conditional convergence cases, the sum of the expected values of the subgames, which is just $\sum_{n=1}^{\infty} 1/2^{\pi(n)} = 1$, does not equal the expected value of the sum of the values, which is somewhere between -99 and -100 .

There are now three versions of the paradox depending on how the game is implemented. The first two versions correspond to the two versions of Satan's Apple: there is a synchronic version where the agent chooses a single betting profile for the game as a whole, and a supertask diachronic version where the agent decides on the bets one by one, with a final payoff at the end. But there is also a third version where the bets are made at a non-supertask pace, say one a minute, for an infinite amount of time.

4.2 *Synchronic version

On the synchronic version of the game, the agent chooses a subset of the subgames to play. For any finite set A of subgames to play, her expected payoff will be $\sum_{n \in A} 1/2^{\pi(n)}$ dollars, namely precisely the amount that she is paid to play the subgames, since the expected value of playing each individual subgame is exactly zero.

Thus, the larger the set A is, the better for the agent. And including one more subgame in the subset will always positively impact the expected payoff. But including *all* the subgames will result in a certainty of losing at least \$98 on balance. The

structure here is very similar to the synchronic version of Satan's Apple. The main difference is that the certainty of profiting from each slice is replaced with an expected profit from it.

Just as in the synchronic version of Satan's Apple, causal finitism does not appear to help. However, just as in Satan's Apple, one can argue that there is no paradox here, just a case with no optimum.

4.3 **Diachronic infinite future version*

On the diachronic infinite future version, an agent who lives infinitely long is regularly asked whether she wishes to play. Each time she rationally should play. Yet her lifetime overall payoff is guaranteed to be negative.

Causal finitism does nothing to help with this case either, as causal finitism does not rule out such future infinities.

However, there is a value problem with this version of the paradox. While I quantified payoffs in dollars, it can't *just* be a matter of gaining or losing money, since in the diachronic non-supertask version you would never get around to spending the money, as the story lasts forever. A natural way to realize this version of the paradox is to think that you simply enjoy or are pained by the payoffs as they come in, by an amount proportional to the payoff.

However, it is a mistake to think that utilities enjoyed over time in general sum to produce an overall utility. Suppose, for instance, that on day n if n is even you get a pleasure of magnitude $1/n$ and if n is odd you get a pain of magnitude $1/n$. It is tempting but mistaken to say that the total value of all these pleasures and pains is just $\sum_{n=1}^{\infty} (-1)^n/n$. For as long as all we are looking at are the values of the pains and pleasures, and not, say, of memories of them, the order in which you receive the pains and pleasures should be irrelevant. This may become clearer if we suppose—in order to make clearer the effect of memory—that at the end of each day, the memories of the pleasures and pains are wiped, so a run of pleasures doesn't get boring and a run of pains doesn't lead to despair. Then it really does seem that permutation of the days shouldn't change the overall outcome.

But we shouldn't be so quick as there is something very counterintuitive here. If permutations don't change overall outcome in the absence of memory and other order-relevant features, then it is no better overall to receive a fixed pain on days whose number is divisible by four and a fixed pleasure on the other days than the other way around. One might think that surely it is better to receive the pleasure on the days not divisible by four, since then any successive sequence of four days will include three days of pleasure and one of pain rather than reversely. But absent memories and other order-relevant features of life, it is arbitrary whether we divide up life into consecutive quadruples of days or into some other collection of disjoint quadruples. And the scenario where "only" the days divisible by four get the pleasure can also be divided up into quadruples with three days of pleasure and one of pain: 4, 8, 16, 1; 20, 24, 28, 2; 32, 36, 40, 3; 44, 48, 52, 5; . . . I contend that there is no real paradox here.

The reason it seems paradoxical is because ordinarily in life order *does* matter—it is miserable to remember that the last four days were mostly painful.

Now, famously, and crucially for Beam's story, you can rearrange the terms in the sum $\sum_{n=1}^{\infty} (-1)^n / n$ to get any real number you like. Since the setup of the story does not include any order-relevant features like memory (if it did, it would be a relevantly different story), in such a case we cannot identify the overall utility over an infinite future life with a sum taken in some particular order. The infinite future version of Beam's story is, thus, a case where there is no meaningful utility that can be attached by summing up the daily utilities—and perhaps no meaningful overall utility at all. Yet the rationality paradox—the paradox that one would lose overall while doing the rational thing each day—would require an overall utility obtained by summing the daily utilities.

4.4 **Diachronic supertask version*

The diachronic supertask version of the paradox comes in two versions. One version is just like the diachronic non-supertask version in that the payoffs are enjoyed as the subgame is played rather than afterwards. In that case, the answer given in the diachronic non-supertask version applies just as well: the value of the whole sequence of payoffs is not equal to the sum of the values.

But there is a particularly troubling version, which is where there is a guarantee that the value of the whole sequence of payoffs equals the sum of the values, for instance, because the terms of the subgames guarantee you will receive an amount of money equal to the total payoff once the supertask is over. In this version, we have the problem that it is rational to accept each offer, but, just as in the case of Satan's Apple, it is bad to accept all of them.

This particularly troubling version is, however, neatly handled by causal finitism, since it requires the final payment made to the agent to be affected by the outcomes from infinitely many games.

4.5 *Evaluation of Beam's paradox*

There are a number of versions of Beam's paradox. Only one of them is resolved by causal finitism, but it is the particularly problematic one. The other versions can be solved without invoking causal finitism. Again, as in the case of Satan's Apple, we get some evidence for causal finitism, but not as much as we would if we could resolve all the versions uniformly by using causal finitism.

5. Evaluation of Decision-Theoretic Paradoxes

Causal finitism gives a resolution of some very interesting intra- and interpersonal guessing paradoxes, and of the most problematic variants of Satan's Apple and of

Beam's paradox. It does not resolve the other versions of Satan's Apple or of Beam's paradox, nor the "parody" rearrangement puzzle in Section 2.6, but these can be handled in other ways.

We get evidence for causal finitism, though not as strong as we would if we could resolve all the variants by using causal finitism. However, because it is the most troublesome versions of the paradoxes that are resolved by causal finitism, consideration of these paradoxes gives us evidence in favor of causal finitism.

Appendix: *Proof of the Theorem from Section 2.1

Let Ω be the space of backwards-infinite sequences of zeroes and ones which we will write as $c = (\dots, c_{-2}, c_{-1}, c_0)$. Order Ω with right-to-left lexicographic ordering (e.g., $(\dots, 0, 0, 0, 1, 0) < (\dots, 0, 0, 0, 0, 1)$). Let the n_k and f_k be as in the statement of the Theorem.

Fix $k \leq 0$. I claim that there is a smallest (with respect to the right-to-left lexicographic ordering) element c of Ω such that for all $k \leq i \leq 0$ we satisfy the constraint

$$f_i(c_{i-n_i}, \dots, c_{i-1}) \neq c_i.$$

To see this, first note that there is at least one element c satisfying these constraints. We can define such an element by finite recursion. Let $c_i = 0$ for $i < k$. Then for $i \geq k$, if c_{i-1} is defined, let

$$c_i = 1 - f_i(c_{i-n_i}, \dots, c_{i-1}).$$

This defines c in a way that guarantees it satisfies the requisite constraints.

For any $n \leq 0$, let Ω_n be the subset of Ω consisting of sequences a such that $a_i = 0$ for $i < n$ and let $\pi_n : \Omega \rightarrow \Omega_n$ be the obvious projection such that $(\pi_n(a))_i = a_i$ if $i \geq n$ and $(\pi_n(a))_i = 0$ otherwise, for any $a \in \Omega$. For any $n \leq 0$, let

$$n^* = \min\{i - n_i : n \leq i \leq 0\}.$$

It is easy to verify that if c satisfies the requisite constraints, so does $\pi_{k^*}(c)$. Moreover, $\pi_{k^*}(c) \leq c$, so for any member of $\Omega - \Omega_{k^*}$ that satisfies the constraints, there is a smaller member of Ω_{k^*} that does so. Since there are only finitely many (indeed, 2^{k^*+1}) members of Ω_{k^*} , it follows that there is a smallest member of Ω_{k^*} satisfying the requisite constraints, and that member must also be the smallest member of Ω satisfying them.

Denote that smallest member of Ω satisfying the constraints by $c^{(k)}$. Observe that $c^{(0)} \leq c^{(-1)} \leq c^{(-2)} \leq \dots$ (since if c satisfies the constraints for one value of k , it does so also for larger ones). Fix $n \leq 0$. Then π_n preserves non-strict right-to-left lexicographic ordering, so $\pi_n(c^{(0)}) \leq \pi_n(c^{(-1)}) \leq \pi_n(c^{(-2)}) \leq \dots$. Thus, the sequence $\pi_n(c^{(-m)})$, for $m \geq 0$, is a monotone non-decreasing sequence in the finite set Ω_n . Hence the sequence must eventually be constant—for large enough m , we must have $\pi_n(c^{(-m)}) = \pi_n(c^{(-(m+i))})$ for all $i \geq 0$. It follows that the rightmost $n + 1$ elements of $c^{(-m)}$ will be constant for large enough m . Since this is true for arbitrary n , we can now define a limiting sequence by $c_k = \lim_{m \rightarrow \infty} (c^{(-m)})_k$. The pointwise limit is always attained.

Now fix $k \leq 0$. Then if we take $m \geq |k|$ to be large enough, we will have $c_i = c_i^{(-m)}$ whenever $k^* \leq i \leq 0$. Since $c^{(-m)}$ satisfies

$$f_i(c_{i-n_i}^{(-m)}, \dots, c_{i-1}^{(-m)}) \neq c_i^{(-m)}$$

for $-m \leq i \leq 0$, it follows by taking m sufficiently large that:

$$f_i(c_{i-n_i}, \dots, c_{i-1}) \neq c_i$$

for $k \leq i \leq 0$. And this completes the proof.

6

The Axiom of Choice Machine

1. Less Technical Introduction

Consider a set S all of whose members are themselves non-empty sets. For instance, take the set

$$S = \{\{1, 2\}, \{1, 4\}, \{-3, 4, 5, 111\}, \{0\}\}.$$

Then it should be possible to “choose” a member from every set in S . More precisely, there should be a function f defined on S that given a member A of S picks out a member $f(A)$ of A .

For instance, in the above example, we could set $f(\{1, 2\}) = 2$, $f(\{1, 4\}) = 4$, $f(\{-3, 4, 5, 111\}) = 111$, and $f(\{0\}) = 0$. This choice was made by means of a rule: for each member of S choose the largest member of that member. But there are many other ways of choosing a member from each member of S , and some of them do not have a briefly describable rule like the above.

The Axiom of Choice (AC) says that for any set S whose members are non-empty sets, there is a “choice function” that for every member A of S “chooses” a member $f(A)$ of A . This is obviously true in many cases. For instance, if every member A of S is a set of positive integers, we could just let $f(A)$ be the smallest member of A (there are other options, too). If every member A of S is a set of integers, we could just let $f(A)$ be the member of A closest to 0, with ties between $-x$ and x broken in favor of the positive (or negative) number (again, there are many options). If every member S is an interval of the form (a, b) with $a < b$ being finite numbers, then we can let $f((a, b))$ be the midpoint of the interval (a, b) (or a point three quarters to the right, etc.). In these kinds of cases, the existence of the “choice function” can be proved from other axioms of set theory.

Similarly, when the set S is finite, the existence of a choice function can be proved from other axioms of set theory.¹ The difficulty is when S is infinite and there is no “natural” way to specify a member of every member of S , as there was in the above numerical examples.

Intuitively, when S is infinite, there should be *more* choice functions. In the case of the set that I began the chapter with, the number of choice functions is $2 \cdot 2 \cdot 4 \cdot 1 = 16$

¹ *One uses mathematical induction on the number of members of S .

(two choices in the first-listed member, i.e., in $\{1, 2\}$; two choices in the second; four in the third; and only one option in the last). If we were to add more members to S , we would intuitively only increase the number of choice functions, by giving ourselves more options.

Nonetheless, this intuition goes beyond the other axioms of Zermelo–Fraenkel set theory. Paul Cohen famously showed that if the other axioms of set theory are consistent, then AC cannot be proved from them, while earlier Kurt Gödel showed that if the other axioms are consistent, then AC cannot be *disproved* on their basis (see Jech 1973).²

There are, however, two kinds of reasons to be suspicious of AC. First, one might have a philosophical view of functions as something like *rules* for mapping, and hence think that functions should be in some sense describable or constructible. If one has such a view, then in the cases where we can give an explicit rule for the choice function—as in the example where we have a set of sets of positive integers and we specify that we always choose the smallest member—one will be happy to admit that there is a choice function. But where no such description can be given, one will think there is no choice function.

Second, there are some paradoxes that can be proved using AC. The most famous of these is the Banach–Tarski paradox (Wagon 1994), which says that a solid mathematical ball can be decomposed into five subsets which can be shifted and rotated to make *two* solid balls of the same radius as the original. In Chapter 5 we considered another paradox, namely that by employing AC we can come up with a guessing procedure that leverages infinitely many past die rolls to make the number of errors finite.

Most working mathematicians put such worries aside and are happy to use AC in their mathematics, and many pivotal and uncontroversial mathematical theorems outside of set theory are proved using AC.

In this chapter, I will give a more precise exposition of some of the mathematical paradoxes of AC, and show that they depend on a weaker version of AC, the Axiom of Choice for Countable Sets of Reals (ACCR). Next, I will argue that ACCR is actually true. The argument will depend on the possibility of an infinite multiverse, which violates finitism but is compatible with causal finitism, as well as on considerations of causally independent stochastic processes. I will then argue that if one can causally compute the values of a choice function through an “ACCR Machine”, then we can make compelling rationality paradoxes out of the paradoxes of AC. For instance, I will show that in coin-tossing scenarios, given plausible symmetry assumptions, one can generate a Dutch Book using ACCR by reasoning parallel to the Banach–Tarski paradox.

Finally, I will argue that if causal infinitism is true, one should be able to construct an ACCR Machine for each situation where one needs to compute a choice function

² *If the axioms are inconsistent, then AC can be *both* proved and disproved from them. And by Gödel’s incompleteness theorem, if they are consistent, then they cannot be proved to be consistent.

for one of our paradoxes. Hence, we have an argument of familiar form: If causal infinitism holds, then there can be ACCR Machines; if there can be ACCR Machines, multiple paradoxes result; so, causal infinitism is false. Also, the ACCR Machine fills a gap in our construction of a countably infinite fair lottery in Chapter 4, Section 3.4 and the die-guessing paradox in Chapter 5, Section 2.

There is a structural similarity between the position I defend on ACCR and the position I defend in the book *vis-à-vis* infinitism. Infinitism is true: it is possible to have an infinite number of objects. But *causal* infinitism is false. Similarly, ACCR (and quite likely full AC) is true but it cannot be made causal: an ACCR Machine is impossible.

Now on to the technicalities, which the non-technically minded reader can skip.

2. *The Axiom of Choice for Countable Collections of Reals

The version of AC that I will be interested in says that if S is a set of non-empty pairwise-disjoint sets (i.e., if $A \neq B$ are members of S , then $A \cap B = \emptyset$) such that each member of S is a countable set of real numbers, then S has a choice function. I will call this “ACCR” (AC for Countable Collections of Reals).

ACCR includes several restrictions over full AC. First, each member of S must be countable: no choices between uncountably many alternatives are required. Second, each member of S is a set of real numbers. Third, the sets are pairwise-disjoint.

If ACCR is true, then any set S that satisfies the conditions of ACCR has cardinality at most \mathfrak{c} , the cardinality of the continuum, i.e., $\mathfrak{c} = \|\mathbb{R}\|$. For if ACCR is true, then a choice function for S provides a one-to-one map from S to (a subset of) the reals. This fact will be important for the construction of a Choice Machine, or, more precisely, an ACCR Machine.

One gets an equivalent reformulation of ACCR if one replaces real numbers in the statement of ACCR with members of any other set that can be proved in ZF (without Choice, of course) to have the same cardinality as the set of the reals. For if we always have a choice function for non-empty pairwise-disjoint countable subsets of \mathbb{R} , then given a collection A of non-empty pairwise-disjoint countable subsets of S , where $\|S\| = \|\mathbb{R}\|$, we can let ψ be a one-to-one function from S onto \mathbb{R} (“onto” means that every element of \mathbb{R} is the result of applying ψ to some member of S), and then consider the non-empty pairwise-disjoint collection of countable subsets of \mathbb{R} given by

$$A' = \{\psi[U] : U \in A\},$$

where $\psi[U] = \{\psi(x) : x \in U\}$. By ACCR there is a choice function f' for A' . It is easy to check we can then define a choice function f for A by letting $f(U) = f'(\psi[U])$. Hence the Axiom of Choice for non-empty pairwise-disjoint countable subsets of \mathbb{R} implies one for such subsets of S ; the converse is proved similarly.

For instance, one gets an equivalent formulation if one replaces the set of real numbers with any fixed non-degenerate interval of real numbers, say $[0, 1]$, since any non-degenerate interval can be proved to have the same cardinality as the set of reals.³

Likewise, one can replace real numbers by members of the set 2^ω of all countably infinite zero-one (or, if one prefers, heads/tails) sequences, since as is well known $\|2^\omega\| = \mathfrak{c}$,⁴ as well as by the members of the set $\{1, 2, 3, 4, 5, 6\}^\omega$ of all countably infinite outcomes of a die roll as that set has the same cardinality (one can encode a die roll as a sequence of three binary digits).

I will now sketch three known paradoxes of AC and observe that they only require ACCR.

3. *Paradoxes of ACCR

3.1 *Die-guessing games*

In Chapter 4, Section 2.4, I gave a paradoxical construction going back to an idea of Gabay and O'Connor showing that given a backwards-(and only backwards-) infinite sequence of die rolls, one can manufacture a strategy that is guaranteed to allow one to guess die rolls with only finitely many mistakes.

The strategy applies the Axiom of Choice to a certain collection of sets of countably infinite sequences of die rolls. These sets are equivalence classes under the relation \sim of differing in only finitely many places. Each equivalence class has only countably many members, since given a single sequence, there are only countably many sequences that differ from it in a finite number of places.⁵ Hence the paradox embodied in the strategy only involves ACCR.

3.2 *Non-measurable sets*

Famously, given AC it is possible to prove that there is a non-measurable subset of the real line, a subset that has no definable “length”. A function μ is a measure on a set U

³ For instance, the function $\tan(\pi(x - 1/2))$ is a one-to-one function from $(0, 1)$ onto \mathbb{R} . Thus $\|(0, 1)\| = \|\mathbb{R}\|$. Hence $\|\mathbb{R}\| = \|(0, 1)\| \leq \|[0, 1]\| \leq \|\mathbb{R}\|$. It follows from the Schröder–Bernstein Theorem (Lang 2002, p. 885), which does not require Choice, that $\|\mathbb{R}\| = \|[0, 1]\|$.

⁴ Any real number in $[0, 1]$ can be written uniquely as $0.x_1x_2x_3\ldots$ in the decimal system, subject to the convention that we prefer an infinite string of trailing zeroes to an infinite string of trailing nines. Thus, any real number defines a unique member of 2^ω , namely the sequence whose n th element is 1 if x_n is odd and is 0 if x_n is even. Hence, $\|[0, 1]\| \leq \|2^\omega\|$. Conversely, any sequence x_1, x_2, \ldots of zeroes and ones defines a different decimal number $0.x_1x_2x_3\ldots$ (the only time two different digit sequences can define the same decimal number is if one of them ends with a string of trailing nines, and that doesn't happen here), so $\|2^\omega\| \leq \|[0, 1]\|$. By Schröder–Bernstein, $\|2^\omega\| = \|[0, 1]\| = \mathfrak{c}$.

⁵ We can show that the equivalence class $[\alpha]$ of a sequence $\alpha = (\ldots, a_{-2}, a_{-1}, a_0)$ has the same cardinality as \mathbb{N} as follows. For any two numbers a and b in $\{1, 2, 3, 4, 5, 6\}$, let $\phi(a, b)$ be the unique member of $\{0, 1, 2, 3, 4, 5\}$ equal to $a - b$ modulo six. For any $\beta = (\ldots, b_{-2}, b_{-1}, b_0) \in \alpha$, let β^* be the number whose base six representation is $\ldots \phi(a_{-2}, b_{-2})\phi(a_{-1}, b_{-1})\phi(a_0, b_0)$ (there will be infinitely many zeroes up front). Conversely, for any natural number n , there will be a unique member β of $[\alpha]$ such that $\beta^* = n$. Thus, we have a bijection between $[\alpha]$ and \mathbb{N} .

provided that it is defined on a σ -algebra of subsets of U (i.e., a set \mathcal{F} of subsets of U that is closed under complements and countable unions) and satisfies the axioms:

- (i) $0 \leq \mu(A)$
- (ii) $\mu(A_1 \cup A_2 \cup \dots) = \mu(A_1) + \mu(A_2) + \dots$ for any countable pairwise-disjoint sequence A_1, A_2, \dots

For a measure μ on the reals \mathbb{R} to be a “length”, we need two further constraints. First, μ is defined on all intervals and $\mu([a, b]) = b - a$ if $a \leq b$, and second, we have translation invariance: $\mu(x + A) = \mu(A)$ for any set A for which μ is defined and any real x , where $x + A = \{x + y : y \in A\}$.

It is easy to check that the standard proof by which non-measurable sets are constructed as Vitali sets (Rudin 1987, pp. 53–4) only applies AC to a set that satisfies the conditions for ACCR. That proof shows that there is a subset of \mathbb{R} , and even a subset of $[0, 1]$, that has no length, if length is assumed to satisfy the above conditions.

This is in itself somewhat paradoxical. It also has the problematic consequence that if a dart is uniformly randomly shot at an interval $[0, 1]$, then there is a subset of that interval such that there is no meaningful probability that the dart will land in that subset, if we require probability to be a countably additive measure, as is standardly done (and with good reason: see Easwaran 2013).

3.3 Banach–Tarski paradox

Subsets of the reals with no length are bad enough. But there are more paradoxical things yet. The Banach–Tarski paradox states that if B is a (solid) ball in three-dimensional space \mathbb{R}^3 , then B can be decomposed into five subsets that can be reassembled into two balls of the same size as B . More precisely, there are five pairwise-disjoint sets B_1, \dots, B_5 of points in \mathbb{R}^3 such that $B = B_1 \cup B_2 \cup B_3 \cup B_4 \cup B_5$, and rigid motions (combinations of rotations and translations) g_2, g_3, g_4 , and g_5 such that B equals $B_1 \cup g_2 B_2$, with B_1 and $g_2 B_2$ disjoint, and $g_3 B_3 \cup g_4 B_4 \cup g_5 B_5$ is a disjoint copy of B , with $g_3 B_3, g_4 B_4$, and $g_5 B_5$ also disjoint.⁶

Here it would be nice to include a picture of the decomposition. But it can’t be done. The decomposition necessarily has to involve “messy” non-measurable sets that cannot be drawn. A decomposition into measurable sets would have to preserve total volume, and it is precisely the point of the paradox that one can double the volume.

The proof of the Banach–Tarski paradox uses AC. But if we examine the proof in Wagon (1994) carefully, we can see that the proof only needs ACCR, as it only applies AC to a certain collection of disjoint countable subsets of a ball, and a ball has the same cardinality as the real numbers.⁷

⁶ For a history and a popular account, see Wapner (2007). For an excellent account of the mathematics of the paradox and related issues, see Wagon (1994).

⁷ Clearly $c \leq \|B\|$ if B is a ball, since a ball contains a line segment. Moreover, $\|B\| \leq \|\mathbb{R}^3\|$. Then, by Cantor–Schröder–Bernstein, it’s enough to show that \mathbb{R}^3 has cardinality c . But \mathbb{R}^3 has the same cardinality

4. *An Argument for ACCR

Despite the paradoxes, there is good reason to think ACCR is true.

What I need to argue is that if S is a set of pairwise-disjoint countable sets of numbers from $(0, 1)$ (which has the same cardinality as \mathbb{R}), then there is a function f on S such that $f(A) \in A$ for all $A \in S$. To avoid triviality, suppose $S \neq \emptyset$.

Imagine now a multiverse consisting of island universes, one per member of S . Suppose that a ψ -particle is a particle with the following property. Once it comes into existence, it lives for a random amount of time whose length is uniformly distributed over the interval $(0, 1)$. Just as it perishes, it spawns a new ψ -particle (perhaps distinguished from the parent particle by a different charge or spin) with its own independent random lifespan. Further, in each island universe, a single ψ -particle comes into existence at the very beginning of time, and each island universe has an infinite future time sequence. There is at any one time at most one ψ -particle in any island universe, and any ψ -particles were generated from a sequence of ψ -particles originating from the first one. Finally, the island universes are isolated from one another in such a way that random events in each universe are independent of the events in the other universes.

For any island universe u , let $\ell_\psi(u)$ be the set of lengths of lives of the ψ -particles in u . For a set A of numbers in $(0, 1)$, we say that an island universe u *matches* A provided that $\ell_\psi(u) = A$.

I now claim that given the above assumptions:

- (1) The following is metaphysically possible: For each member A of S there exists exactly one matching island universe u .

Given (1), there is an argument for the existence of a choice function. First, note that it is widely held that purely mathematical objects cannot exist merely contingently, and choice functions for collections of sets of real numbers are purely mathematical.⁸ Thus, if we could show that it is possible that a choice function exists for S , it will follow that a choice function *actually* exists for S .

Now, suppose that we are in a multiverse satisfying the condition in (1). Then for any $A \in S$, we can define $f(A)$ as follows. There is a unique matching island universe u . Then let $f(A)$ be the length of life of the first ψ -particle in u . Then $f(A) \in \ell_\psi(u) = A$, and so we have a choice function. Hence, if the scenario that (1) claims to be possible were actual, there would be a choice function. Since the scenario is possible, there actually is a choice function.

as $(2^\omega)^3$, as \mathbb{R} has the same cardinality as 2^ω . But there is an easy one-to-one map from $(2^\omega)^3$ onto all of 2^ω : just let $f((a_1, a_2, \dots), (b_1, b_2, \dots), (c_1, c_2, \dots)) = (a_1, b_1, c_1, a_2, b_2, c_2, \dots)$.

⁸ It is also widely thought by philosophers that there are sets that are not purely mathematical and that have contingent beings as members. Such sets may exist in some worlds but not others.

Thus, in order to argue for ACCR, it remains to argue for (1). To do that, I want to draw on the following principle:

- (2) Suppose that the P_i are causally independent physical stochastic processes for each value of i in some mathematical index set I , and each P_i generates an output that can be quantified as a mathematical object (say, a sequence of numbers quantifying some physical quantity). Let “ X_i ” be an abbreviation for “the mathematical quantification of the output of P_i .” Suppose that for each i , $Q_i(x)$ is a mathematical formula with no free variables besides i and x . Then if we have $\forall i \in I (\Diamond Q_i(X_i))$, we also have $\Diamond \forall i \in I (Q_i(X_i))$,

where the diamond indicates *causal* possibility, and $\forall x \in A (F(x))$ abbreviates $\forall x (x \in A \rightarrow F(x))$.

Normally, to interchange the order of a possibility operator and a universal quantifier is fallacious. For instance, every natural number n is such that possibly there are exactly n horses in existence, but we cannot infer from this that possibly every natural number n is such that there are exactly n horses in existence.

But what (2) claims is that we can interchange a causal possibility operator and a universal quantifier when we are making claims about the outcomes of independent physical stochastic processes. For instance, if a countably infinite number of dice are tossed, then for each n , it is causally possible that the n th die shows the number 3. And likewise it is causally possible that all the dice show the number 3. Otherwise, intuitively, there would have to be some kind of dependence or coordination between the dice ruling out every die’s showing the number 3. Similarly, for all n it is possible that the n th die shows one plus the remainder after dividing n by six, and likewise, by independence, it is causally possible that for all n , the n th die shows one plus the remainder after dividing n by six (i.e., that the sequence of dice yields 2, 3, 4, 5, 0, 1, 2, 3, 4, 5, 0, ...).

Note that the restriction that $Q_i(x)$ is a *mathematical* formula keeps $Q_i(x)$ from making physical claims about the outcomes of the other physical processes P_j , which would allow for counterexamples.

The sense of independence in this argument is a causal sense. A paradigm example of independence of this sort is where the stochastic processes occur in causal isolation from each other, say in causally isolated island universes.

Normally, in probability theory we say that a finite collection of random variables X_1, \dots, X_n is independent if and only if for every sequence A_1, \dots, A_n of measurable sets,

$$P(X_1 \in A_1 \ \& \ \dots \ \& \ X_n \in A_n) = P(X_1 \in A_1) \cdot \dots \cdot P(X_n \in A_n).$$

And then we say that an infinite collection of random variables is independent if and only if every finite subcollection of it is independent.

The probabilistic characterization is known to be insufficient to characterize *genuine* independence. Fitelson and Hájek (2017) argue that this characterization fails

when we are dealing with zero probability events. Here is an example of such a failure. I simultaneously toss an infinite number of coins. These coins behave like independent fair coins do, with one exception. If all the coins but the first (in some numbering system) land heads, they simultaneously cause the first coin to land tails. (Yes, this violates causal finitism, but it's just an intuition pump.) The only thing that disturbs independence is this extra causal intervention. But this intervention only occurs if the infinitely many tosses after the first are all going to be heads, and the probability of that happening is zero. So the probability of the disturbance is zero, and modifications to processes that only have a probability of zero of occurring do not affect the values of any unconditional probabilities, and hence do not affect probabilistic independence. However, even though we have probabilistic independence, we clearly do not have causal independence.

In this example, we would have a violation of the conclusion of (2). Let the two outcomes of a coin toss be quantified as 0 (tails) or 1 (heads), and let $Q_i(x)$ say that $x = 1$. Then for all i , it is causally possible that $Q_i(X_i)$, i.e., that the i th toss is heads, but it is not causally possible that for all i we have $Q_i(X_i)$, i.e., that all the coins land heads. And it is precisely the violation of causal independence which resulted in the violation of (2). That is why in (2) we need causal independence and not just probabilistic independence.

Now, given (2), we get (1). For let g be a one-to-one function from S to all the island universes. Let $I = S$. Then for $A \in S$, let X_A be the sequence of ψ -particle lifetimes in the world $f(A)$. Let $Q_A(x)$ say that x is a sequence such that the set of its elements equals A . Thus, $Q_A((x_1, x_2, \dots))$ says that $\{x_1, x_2, \dots\} = A$.

Now, any countable set of numbers in the interval $(0, 1)$ can be the set of lifetimes of ψ -particles in the island universe $g(A)$. Thus, $Q_A(X_A)$ is causally possible for each $A \in S$. Hence by (2), it is causally possible that for all $A \in S$ we have $Q_A(X_A)$. But what is causally possible is also metaphysically possible. So it is metaphysically possible that for all $A \in S$ we have $Q_A(X_A)$. But to say that $Q_A(X_A)$ is another way of saying that island universe $f(A)$ matches A , and so we have (1) as desired, which completes our argument for ACCR.

The above argument presupposes that sets built out of real numbers cannot exist merely contingently, and hence if there *could* exist a choice function (functions are just sets of ordered pairs), then there actually *does* exist a choice function. One might question this assumption. For instance, perhaps the correct metaphysics of sets is an Aristotelian rather than Platonic one, and hence the only sets that exist are the ones that can be obtained by abstraction from the actually existing things (Pruss ms). On such a view, in the world where (1) holds, a choice function exists—but it might not exist in the actual world.

If this is so, then the argument for ACCR fails. But the paradoxes we will discuss involving ACCR could be run in a possible world where the requisite choice function

exists, and would be hardly less paradoxical. Thus, even if this argument for ACCR fails in this way, this is of no help in escaping paradox.

5. *A Choice Machine

5.1 *Strange mathematics and paradox*

When the Banach–Tarski paradox was explained to him, the celebrated physicist Feynman (1985, p. 85) said he wasn't bothered because the paradox couldn't work for a real orange. The paradox essentially depends on a ball that is a continuum. Oranges, however, are made of a finite number of discrete particles, and hence do not admit of a paradoxical disassembly. Analogously, our other paradoxes of Choice aren't realized in our physical reality.

However, the paradoxes become compelling if they present problems for rationality, even when the situations they concern are not physically realizable given the actual constitution of our world. Rationality as such should not be tied to a particular world (cf. Chapter 4, Section 4.3). I have already argued in Chapter 4 that the die-guessing strategies of Gabay and O'Connor yield a paradox for rationality. What about the other mathematical paradoxes?

We in fact don't have any chance to bet on non-measurable sets of dart throws. But we might think we do, and there should be a sensible answer to the question of how we should act when we do, *if* such situations are possible. If such situations are not possible, on the other hand, then it should not surprise us if we get answers that don't seem sensible. After all, we should in general not expect a sensible answer to the question of how you should bet if you believe something metaphysically impossible. How, after all, should you bet on a coin toss if you think the coin has a chance of 1.7 of landing heads and a chance of -0.3 of landing tails?

The metaphysically possible realizability of betting situations based on the paradoxes is, thus, more of a problem than the purely mathematical puzzles. As purely mathematical puzzles, they teach us only that mathematical reality is strange, which should not come as a great surprise for us.

The die-guessing game and the nonmeasurability problem lend themselves directly to betting scenarios. But it is only if we can have a causal process that embodies a choice function that we can realize the paradoxical strategy in the die-guessing paradox, and it is only if a payoff can be made to depend on whether a dart lands in a non-measurable set that one can have a compelling betting paradox. But to do that, we need something like a Choice Machine.

The Banach–Tarski result is a particularly curious piece of mathematics. To make it compelling as a paradox, we can tie it to a betting scenario. Suppose that a point is randomly and uniformly chosen in a cubical region of space, with four specially

distinguished ball-shaped subregions consisting of balls each of which has volume $1/100$ th of the volume of the whole cube.

A betting portfolio will be a finite list of regions of outcome space (the cubical region of space) together with the payoff if the outcome lies in the given region. Suppose you are a rational agent. You should then be happy to pay a dollar for a betting portfolio where:

- (i) you pay \$1000 if the point is in the first ball-shaped region, and
- (ii) you get \$1200 if the point is in the second ball-shaped region. (Fig. 6.1.)

For your expected payoff is $(1/100) \cdot \$1200 + (1/100) \cdot (-\$1000) - \$1 = \1 .

The following rearrangement principle is very plausible in the case of a uniformly random selection:

- (3) If a rational agent is happy to pay x for a betting portfolio X , and betting portfolio X' differs from X by replacing one of the outcome regions A in X by an outcome region A' that differs from A only by a rigid motion, with the same payoff as A had, then the agent is happy to pay x for X' .

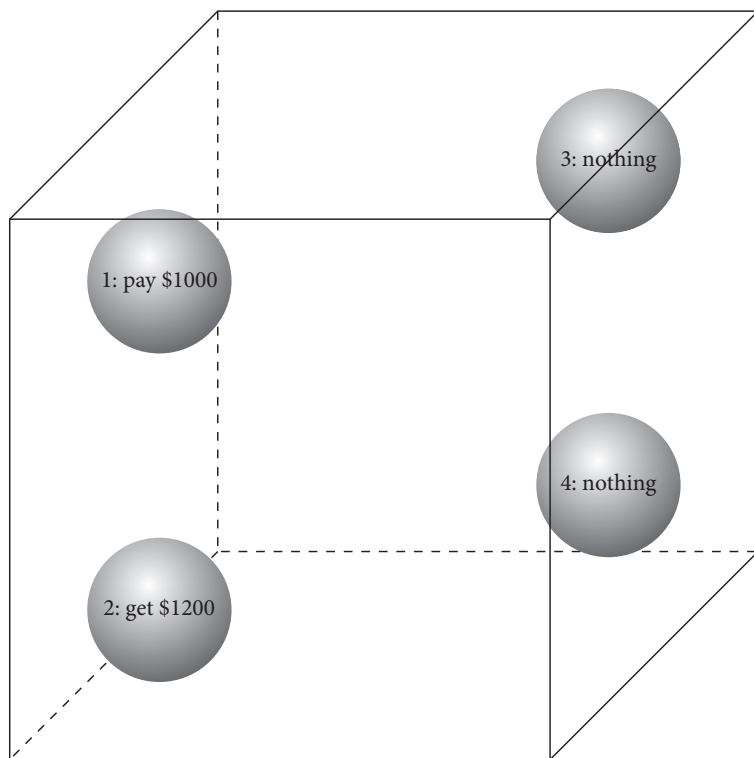


Fig. 6.1 The (i)–(ii) betting portfolio that you should be happy to pay a dollar for. The volume of each sphere is $1/100$ th of that of the cube.

We also have this uncontroversial equivalence principle:

- (4) If a rational agent is happy to pay x for a betting portfolio X , and betting portfolio X' is equivalent to X in the sense that it has the same total payoff at each point in the space of outcomes, then the agent is happy to pay x for X' .

For instance, one way to get an equivalent portfolio is to take some region R in X that has a payoff y , and replace it by a list of regions R_1, \dots, R_n , each with payoff y , where the R_i are disjoint and $R = R_1 \cup \dots \cup R_n$.⁹

Suppose that B_1, \dots, B_5 are five subsets of the first ball-shaped region that can be rearranged into two balls of the same size, as per the Banach–Tarski theorem. By (4), if you were happy to pay \$1 for our initial scenario, you will be happy to pay \$1 for an equivalent portfolio where instead of having to pay \$1000 for an outcome in ball one, you will have to pay \$1000 for an outcome in each of the regions B_1, \dots, B_5 . But now we can form a series of betting portfolios where the regions B_1, \dots, B_5 are rigidly moved, one-by-one, in such a way that at the end they disjointly fill the third

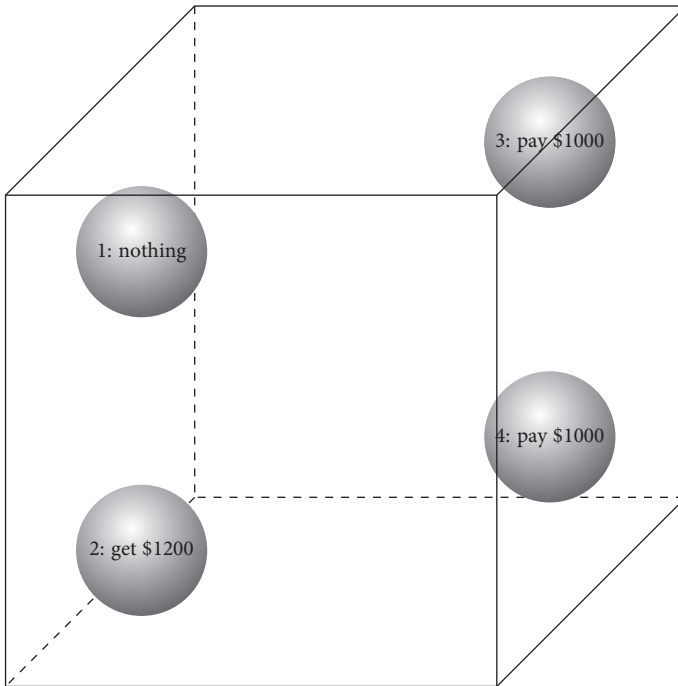


Fig. 6.2 The (i')–(ii') betting portfolio that you should be happy to pay a dollar for if the argument works.

⁹ One might say that the portfolios aren't merely equivalent, but are the same portfolio. However, if we think of portfolios as functions from sets of regions to payoffs, technically the two portfolios are different, even though they prescribe the same respective payoffs under all possible circumstances.

and fourth ball-shaped regions. As long as the payoffs are unchanged, by (3) you will still be willing to pay a dollar for the portfolio.

The result will be equivalent to paying a dollar for this portfolio:

- (i') you pay \$1000 if the point is in the third or fourth ball-shaped region,
- (ii') you get \$1200 if the point is in the second ball-shaped region. (Fig. 6.2.)

So by (4), you will be happy to pay a dollar for this portfolio.

This means that you are happy to pay \$1 for a portfolio whose expected payoff is $(1/100) \cdot \$1200 + (2/100) \cdot -\$1000 = -\$8$. That is absurd.

If one wishes, one can manufacture a Dutch Book using this method, a set of betting portfolios each of which you would be rational in accepting but where you are certain to lose each time. For instance, you would surely also accept for free this scenario:

- (i'') you get \$900 if the point is in the first, third, or fourth ball-shaped region,
- (ii'') you pay \$2500 if the point is in the second ball-shaped region. (Fig. 6.3.)

For the payoff for this scenario is $(3/100) \cdot \$900 + (1/100) \cdot -\$2500 = \$2$. But now we have three scenarios that you would rationally accept: (i)–(ii) for a dollar, (i')–(ii') for a dollar, and (i'')–(ii'') for free. If you accepted all three, then in the first ball-shaped region, you would get $\$900 - \$1000 - \$2 = -\102 , in the second $\$1200 + \$1200 - \$2500 - \$2 = -\$102$, in the third and fourth $\$900 - \$1000 - \$2 = -\102 ,

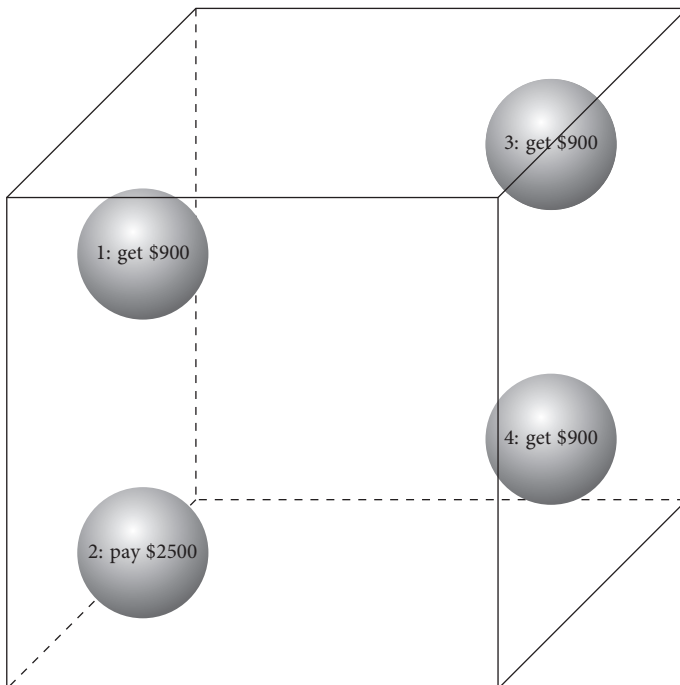


Fig. 6.3 The (i'')–(ii'') betting portfolio you should be happy to accept for free.

and everywhere else you would pay two dollars with no prize. As a result, no matter what, you'd be losing at least two dollars.

If we could implement real betting scenarios based on our strange mathematical stories, we would indeed get paradoxes for rationality. But this requires a way of actually implementing payoffs based on whether results fall in a set “generated” with the Axiom of Choice. And it seems we need a Choice Machine to implement such payoffs.

5.2 Coin-flips and Dutch Books

Perhaps you are not impressed that much by the betting version of the Banach–Tarski paradox, because you do not think there are ways of genuinely uniformly choosing points in three-dimensional space in a way that is invariant under rigid motions.

If so, I can also give a paradox involving coin-flips. Consider the space Ω of countably infinite sequences of heads and tails, corresponding to the results of tossing an infinite number of independent, fair, and indeterministic coins. The coins are all alike, so our reasoning about the coin tosses should be invariant under permutations of coins. More precisely, suppose π is any permutation of the natural numbers \mathbb{N} , and define $\pi^*((\alpha_0, \alpha_1, \dots)) = (\alpha_{\pi(0)}, \alpha_{\pi(1)}, \dots)$. For a subset A of Ω , as usual let $\pi^*A = \{\pi^*(\alpha) : \alpha \in A\}$. We should thus be indifferent between betting scenarios equivalent under π^* for any permutation π of \mathbb{N} .¹⁰ The invariance of what it is rational to predict about coin toss results under permutations of coins is central to the intuitive idea of the coin tosses as independent and fair.

But it turns out that Ω has something like a paradoxical decomposition for mathematical reasons that are closely related to the Banach–Tarski paradox. To be more precise, there is a subset D of Ω with the property that $P(D) = 0$, i.e., there is zero probability¹¹ that the sequence of coin tosses lands inside D , and disjoint subsets A_1, A_2, A_3, A_4 of $\Omega - D$ (where $A - B = \{x \in A : x \notin B\}$) such that

$$\Omega - D = A_1 \cup A_2 \cup A_3 \cup A_4$$

and there exist permutations ρ and τ of the natural numbers such that:

$$\rho^*A_2 = A_2 \cup A_3 \cup A_4$$

and

$$\tau^*A_4 = A_1 \cup A_2 \cup A_4$$

The proof uses ACCR, and will be sketched in the Appendix to this chapter.

¹⁰ Or perhaps we can have infinitesimal preferences between them—i.e., if we're dealing with bounded finite payoffs, we might be willing to pay infinitesimally more for one over another. For instance, maybe we should think it is infinitesimally less likely that all the even numbered coins landed heads than if all the prime-numbered coins landed heads (though see Williamson 2007), even though there is a permutation π that will map one outcome set to the other. However, such infinitesimal preferences won't affect the Dutch Book I will construct, as it rests on non-infinitesimal differences.

¹¹ If one prefers a non-classical setting for probability theory where all options get non-zero probability, we can take $P(D)$ to be infinitesimal, and make minor revisions to the rest of the argument.

It follows from the above that:

$$\Omega - D = A_1 \cup \rho^* A_2$$

and

$$\Omega - D = A_3 \cup \tau^* A_4.$$

Now consider this very plausible betting principle for games based around a countably infinite sequence of independent, fair, and indeterministic coin tosses, a principle analogous to (3):

- (5) If a rational agent is happy to pay x for a betting portfolio X , and betting portfolio X' differs from X by replacing one of the outcome regions A in X by an outcome region $\pi^* A$ for some permutation π , with the same payoff as A had, then the agent is happy to pay x for X' .¹²

Permuting the coins by π shouldn't make a difference to a betting portfolio. In the case of permutations that move around a finite number of coins, (5) corresponds to the property of exchangeability of independent and identically distributed random variables. For instance, given a collection of ten coins, the probability that there will be exactly three heads among coins 1, 2, 3, and 4 is the same as the probability that there will be exactly three heads among coins 4, 5, 6, and 7, there being a permutation of the coins that swaps 1, 2, 3, and 4 with 4, 5, 6, and 7, respectively. In (5) this is extended to infinite permutations, and is formulated in terms of betting rather than probability.

Now, a rational agent would pay \$1.00 for a portfolio where she wins \$1.25 if the sequence of results falls in $\Omega - D$. After all, the probability that the sequence doesn't fall in $\Omega - D$ is zero, and so almost surely she will get \$0.25 per game.

Applying (5) and the decomposition $\Omega - D = A_1 \cup \rho^* A_2$ (let $\pi = \rho^{-1}$ and $A = \rho^* A_2$ in the context of (5)), we conclude that the agent would also pay \$1.00 for a portfolio where she wins \$1.25 if the results fall in $A_1 \cup A_2$. And applying (5) and the decomposition $\Omega - D = A_3 \cup \tau^* A_4$ (let $\pi = \tau^{-1}$ and $A = \tau^* A_4$) lets us conclude that she would pay \$1.00 to play a game where she wins \$1.25 if the results fall in $A_3 \cup A_4$. But if she plays both of these games, then she will pay \$2.00 per game. If the tosses fall in $A_1 \cup A_2 \cup A_3 \cup A_4$, she will win \$1.25; if they fall in D , she will win nothing. So on balance she will lose at least \$0.75 per game. She is thus subject to a Dutch Book.

5.3 How to construct a Choice Machine

5.3.1 ANGELS

The simplest method of generating a Choice Machine is to posit an angel in the setting of the multiverse described in Section 4, a being not limited by physics who observes all the island universes—here's where causal infinitism will be involved—and can read off values of the choice function from the distributions of ψ -particle lifetimes in that universe.

¹² We can further stipulate, if we wish, that the happiness in both cases is non-infinitesimal—i.e., that the agent takes the scenarios to have non-infinitesimally positive value. Cf. note 10, above.

I have supposed such beings in previous arguments, but in those arguments it was easier to see how to mechanize the situation. For instance, in Chapter 5, I considered angels that can announce whether an infinite number of some type of event occurred. It is not that hard to imagine a mechanization of that. For instance, we could suppose a detector for that type of event which shifts an indicator needle by sixteen degrees the first time it detects the event, by eight degrees the next time, by four the next time, and so on. (This may be done in a supertask as needed.) Assuming a continuous physics, infinitely many events of the given type have occurred if and only if the indicator needle has moved by a full thirty-two degrees ($32 = 16 + 8 + 4 + 2 + 1 + 1/2 + \dots$).

It will be a useful exercise to see if one could imagine something that could count as a physical process—even if subject to laws of nature other than ours—which implements a choice function. If we can do that, our ACCR-based arguments from causal infinitism to paradoxes are strengthened.

5.3.2 A FOUR-DIMENSIONAL MACHINE

5.3.2.1 Making the machine Surely, a three-dimensional space is possible. But one version of string theory posits an eleven-dimensional spacetime: ten spatial dimensions and one temporal dimension. It would be surprising if a three-dimensional space were possible but a ten-dimensional one were not, and conversely if a ten-dimensional one were possible, but a three-dimensional one were not. More generally, I take it that we have good reason to think that for every positive integer n , it is possible for space to have dimension n . In particular, it is possible to have four-dimensional space.

There is one little hangup here. Perhaps the property of spatiality is tightly connected to the laws of nature we actually have. Einstein (2015, p. 176), for instance, speculated that spacetime is “a structural quality” of the gravitational field. But, plausibly, a *gravitational* field is something that couldn’t exist with laws of nature other than ours, just as it is essential to water that a water molecule is composed of exactly three atoms. Of course, it is metaphysically possible to have water*, a substance different from water with very similar macroscopic behavior but a different number of atoms per molecule, and it would be possible to have a gravitational* field which behaves similarly to a gravitational field but is not a gravitational field. If this is right, and if spacetime is identical with the gravitational field, and if space is necessarily bound up with spacetime, then it might not be possible to have a different dimensionality of space than we do. But in that case, it would surely be possible to have “space*”, something with a different dimensionality from space but analogous to space in its role in physical reality.

To avoid clutter and because I am not persuaded by the above suggestion about space having essentially the number of dimensions it does,¹³ I will omit the asterisk and simply say that four-dimensional space is possible.

¹³ I prefer the view that space is understood functionally, and a world with other laws but where some determinables play a role sufficiently similar to that played by location in our world is a world with space.

Assume causal infinitism, and work in a world with a continuous five-dimensional spacetime, i.e., four spatial dimensions and one temporal dimension, and all the dimensions quantifiable by means of real numbers. Suppose that S is a collection of pairwise-disjoint countable sets of real numbers in the interval $(0, 1)$; all the other choice situations can be reduced to this by an appropriate bijection. In Section 4 it was already argued that ACCR is true. This has a useful consequence: the cardinality of S is less than or equal to that of the continuum. For there exists a choice function for S , and this will be a one-to-one function from S into $(0, 1)$.

If the cardinality of S is less than or equal to that of the continuum, we can fit into our universe a three-dimensional machine for each member of S , just by putting the machines into different three-dimensional hyperplanes in the four-dimensional space.

Imagine a bunch of infinite three-dimensional machines, which will be parts of the full four-dimensional machine, with one three-dimensional machine for each member of S . I will call each of these machines a “slice”. Each slice then consists of an infinite string of finite sub-machines, which I will call “blocks”, arranged in a sequence with a first block (the “lead block” I will call it), a second, and so on (Fig. 6.4). The blocks each have a knob where any number in $(0, 1)$ can be set. The blocks also each have two wired inputs, α and β , and two wired outputs, γ and δ , each capable of receiving or transmitting an encoding of a number in $(0, 1)$, say, an encoding in some analogue of an electric pulse of a length proportional to that number. Finally, each block has a wireless output capable of omnidirectionally, along all four dimensions, transmitting an encoding of a number in $(0, 1)$, say, by some analogue of an electromagnetic pulse.

Each block then has the following functional properties:

- (i) When a signal of value a comes in via wired input α , if a equals the knob setting on the block, then the signal is relayed through output δ ; otherwise, it is relayed via γ .
- (ii) When a signal comes in via wired input β , then a wireless signal of value equal to the block’s knob setting is emitted.

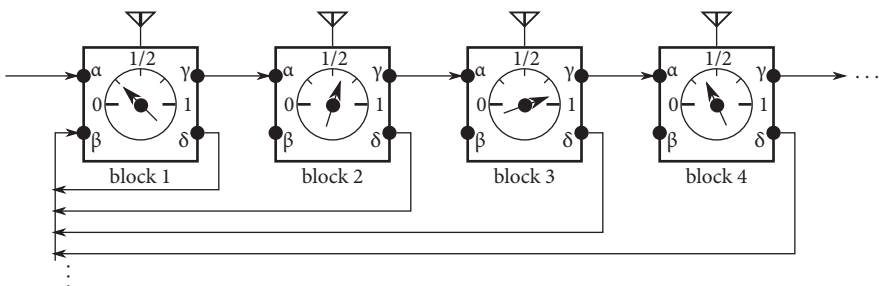


Fig. 6.4 A slice of a Choice Machine.

Moreover, I suppose that wired input β is capable of being connected to infinitely many outputs, and when a signal comes from any one of them, the wireless signal is triggered.

The blocks are wired in each slice as follows for each n :

- (iii) Output γ of block n is wired to input α of block $n + 1$.
- (iv) Output δ of block n is wired to input β of block 1.

Then a wire (running out into the fourth dimension) connects the α inputs of block 1 in each of the slices to the global input of the machine.

The four-dimensional machine then works as follows. A global input is sent to each slice's lead block's input α . In each slice, the input is passed on down the sequence of blocks until it finds a matching knob setting, if there is one. If it does, then a signal is sent to the lead block in the slice, which then emits *its* knob setting wirelessly.

We now say that the full four-dimensional machine *matches* the set S that we want a choice function for provided that (a) no two knobs are set to the same value in the machine, and (b) there is a one-to-one map f from S to all of the machine's slices, such that if $A \in S$, then the set of knob values in slice $f(A)$ equals the set A .

The causal-independence reasoning from Section 4 tells us that it's possible to have such a machine matching S . All we need to do is suppose that the knob settings are random and causally independent and that we got lucky enough to have a match. Randomness is not necessary, though (see Section 5.3.5).

5.3.2.2 Using the machine Now, if the machine matches the set S , then it can be used to compute the values of a choice function for S as follows. Given a member A of S , select any value x in A . Send this value into the global input of the Choice Machine. This signal will then be propagated to the α inputs of all the lead blocks, and then onward to the α inputs of further blocks, until it hits a block whose knob is set to x . There is exactly one block with a knob set to x , and when the signal reaches that block's α input, it gets directed to the β input of the lead block in the relevant slice, which then wirelessly outputs its knob value. This knob value is the value of the choice function as applied to the set A containing x . The set of knob values for that slice is going to be equal to A , and the lead block's knob value will thus be a member of A .

One complication is this. The machine may take a very long time to respond, since it may take a long time for the signal to propagate along the blocks of the relevant slice to the block whose knob value matches the input. As a result, the machine will not be usable for guessing sequential die rolls in the Gaby–O'Connor paradox, since no matter how short the transmission delay between blocks, as long as it's non-zero, there is no guarantee that an answer will be obtained before the next roll. We can fix this problem by modifying the paradox to suppose you have an arbitrarily long finite amount of time to guess each roll before the roll is made. Alternately, we can accelerate the machine, for instance by making sure that transmission delay from block n to block $n + 1$ is $1/2^n$ units of time, and that the transmission delay between each block's

δ output and the lead block's β input is the same no matter how far away the blocks are. This requires either that there be no analogue to the speed of light limit that our world has, or that the successive blocks, and the connections between them, shrink exponentially as the block number goes up.

One may also be a little uneasy about the fact that one must *choose* a value x in A to use the machine. Granted, no matter what value one chooses in A , the output of the machine will be the same—the knob value of the lead block in the same slice. But, still, if we're making a Choice Machine, we shouldn't have to do any choosing ourselves.

Fortunately, this doesn't matter for our applications. For in all our applications, what we need is a choice function f for the set of the equivalence classes under an equivalence relation \sim , and what we will really need to compute is just $f([x])$ for a particular x , where $[x]$ is x 's equivalence class. So the x is given.

For instance, in the Gabay–O'Connor guessing paradox, we are given a past sequence of die rolls, say $\dots, a_{-8}, a_{-7}, a_{-6}$, and we need to evaluate the choice function at the equivalence class that the sequence will fall into. Granted, we don't know $a_{-5}, a_{-4}, a_{-3}, a_{-2}, a_{-1}, a_0$, but since sequences differing in a finite number of places are equivalent, all we need to do is to plug the sequence $\dots, a_{-8}, a_{-7}, a_{-6}, 1, 1, 1, 1, 1$ into the input of our Choice Machine. To do that, of course, we'll need to encode sequences of die rolls into numbers in $(0, 1)$ (e.g., encoding them into a sequence of digits after the decimal point). This is likely to need a supertask. And then we'll need to decode the output similarly, by determining the digits of the output value.

The case of non-measurable sets, Banach–Tarski, and my coin-guessing Dutch Book paradox may be more complicated, but all these cases involve in the end having to check if some point falls in a set generated by ACCR.

5.3.2.3 Causal infinitism and verifying the machine's match The use of causal infinitism in our Choice Machine is subtle. While a supertask may be involved in encoding, say, an infinite sequence of coin tosses, once the encoding is done, only a finite number of blocks are involved in the emission of the output. Nonetheless, the non-activation of the infinite number of other blocks' δ outputs is essential to the machine's having a well-defined output. The output's being what it is is thus causally dependent on the other blocks not matching their input to their knob value. This kind of negative dependence with positive initial causes will be discussed some more in Chapter 7, Section 2.5.

But there is a place in which a less subtle use of causal infinitism will occur. For in order to generate rationality-based paradoxes of infinity, it is not enough to have a Choice Machine for S available. The applications require the agent to *know* that they have a machine matching S .¹⁴ If the Choice Machine is generated by random settings of knobs, the knowledge that a Choice Machine matches S will have to depend

¹⁴ I am very grateful to Ian Slorach for pointing this out to me.

on an infinite number of random events. If this dependence is causal—and perhaps apart from controversial cases where the knowledge is mediated by divine testimony (see Chapter 9, Section 3.4), it is hard to see how it could be anything but causal—we have a clear violation of causal finitism.

This raises, however, the question of *how* one could verify the match of a machine to S . For instance, if the machine was generated by a random process, how do we know that the random knob settings worked out so as to make the machine compute a choice function for S ? After all, it is very unlikely that it would so work out.

One could posit an angel that can tell at a glance that the knob settings match S . The angel would need to verify the following features of the situation:

- (i) Each slice is correct: the set of its knob values is a member of S .
- (ii) There are no duplicates: no knob value occurs in two different blocks.
- (iii) The machine is complete: each member x of a member of S is the knob value of some block of some slice.

Could we check this in a more mechanical way, perhaps with more assumptions about what is physically possible in our $(4 + 1)$ -dimensional universe?

Part (i) is the easiest. In all the applications we need, S is the set of equivalence classes under some equivalence relation, and each equivalence class can be enumerated given any member of it—the equivalence classes in all cases are given by equivalence under the action of some countable group (say, the group of changes to a finite number of die roll results). We could thus suppose that there is a machine in the hyperplane of each slice which travels from block to block, verifying that the knob values enumerate a full equivalence class, and doing so in a supertask to ensure a finite finish time. If there is a failure, it explodes the full machine. Thus, if no explosion happens in the requisite amount of time, (i) has been verified.

Parts (ii) and (iii) are harder.

We can, however, suppose that there is a continuum $(K_x)_{x \in (0,1)}$ of kinds of fields (electromagnetic, shmagnetic, phignetic, etc.), each indexed by a different number x in $(0, 1)$, in such a way that any combination of these kinds of fields can be generated: for instance, one can generate precisely a $K_{0.24}$ and a $K_{1/\sqrt{\pi}}$ field at the same time. Moreover, the fields do not interact with one another. We now suppose that each block has a self-test button which makes it omnidirectionally (in all four dimensions) emit a pulse of field K_x , where x is the block's knob value. Moreover, each block has a receiver calibrated so that if it receives a pulse of field K_x that wasn't transmitted by itself, it explodes the whole machine. We now need to suppose that the verifying machine that travels along the slice presses the self-test button on each block. If there is a duplicate of the block, then the machine explodes. Non-explosion then yields (ii).

Next, suppose a second disjoint continuum $(K'_x)_{x \in (0,1)}$ of kinds of fields indexed by the numbers x in $(0, 1)$, also not interacting between one another or with the original K_x fields. Moreover, suppose that a field of kind K'_x corresponding to each member x of a member of S has been initially generated, and that a block with knob value x

has a disposition to neutralize a field of type K'_x in the vicinity of the device. We give enough time for all the neutralizations to take effect, and then a detector will blow up the whole machine if any of the K'_x fields has not been neutralized. Non-explosion yields (iii).

Of course, there is the issue of how one knows that the verification system—all the field emitters and everything—is in place. Perhaps we can suppose that the rational agents we are running the paradoxes for are capable of infallibly creating any precisely specified physical system. Or perhaps here we do need an angel reporting the verification setup, and so we have perhaps not advanced too far from the angelic “machine”.

5.3.3 A THREE-DIMENSIONAL MACHINE

The description above used a four-dimensional machine to make it easier for the reader to visualize the slices as actual machines. But we could imagine a world where complex two-dimensional machines are possible.¹⁵ In such a world, we might even be able to make a two-dimensional machine for each block, and then the slices could be assembled into a three-dimensional machine.

5.3.4 ** IS AC NEEDED?

All of the four problems considered depended on AC. But could one, perhaps, recreate the same problems without dependence on AC? If one could do that, it would be plausible that one could recreate them causally without Choice Machines.

However, there is good reason to think that Choice cannot be eliminated. Solovay (1970) has shown that if ZF set theory is compatible with the existence of inaccessible cardinals, it is also compatible with both the Axiom of Dependent Choice (DC) and the hypothesis that all sets in \mathbb{R}^n (for all n) are Lebesgue measurable. Thus, if ZF is compatible with inaccessible cardinals—and it is widely assumed that it is¹⁶—then something beyond ZF is needed to prove the existence of non-measurable sets.

Moreover, one can use all three scenarios to prove the existence of non-measurable sets, at least assuming DC (which is needed to have Lebesgue measure). The Banach–Tarski paradox shows that not all sets in \mathbb{R}^3 are Lebesgue measurable, since if they were, so would be the sets involved in the paradoxical decomposition, and then the volume of one ball would equal the total volume of two balls of equal size.¹⁷ The coin toss rearrangement scenario of Section 5.2 shows that the measure on Ω corresponding to the coins being independent, fair, and identically distributed does not make all subsets of Ω measurable. But that measure is isomorphic to Lebesgue

¹⁵ Dewdney (2001) gives a compelling fictional account of a complex two-dimensional world.

¹⁶ There is a large body of mathematical theory going beyond ZFC that presupposes an inaccessible cardinal—Grothendieck’s cohomology theory which uses Grothendieck universes that require the existence of an inaccessible cardinal (Artin, Grothendieck, and Verdier 1972).

¹⁷ In fact, in this one case, one can drop the consistency between ZF and inaccessible cardinals from the argument; see Wagon (1994, Theorem 13.2).

measure on $[0, 1]$ (as usual, use the correspondence between sequences of coin tosses and binary expansions of numbers), and so it follows that there is a Lebesgue non-measurable set. Finally, the die-guessing scenarios also generate a non-measurable set. For, given independence, there can be no *measurable* function from past results to future results that does better than chance at guessing future results, and if there are non-measurable functions, there are non-measurable sets.

5.3.5 LUCK

In the above, I suggested that we could simply get lucky so that the ψ -particle lifetimes satisfy (1) or the knobs on slices satisfy the matching condition, which constraints are needed for our setup to be a successful ACCR machine. But one of the useful things about being able to causally generate choice functions was to generate a countably infinite fair lottery using the more complex construction of Chapter 4, Section 3.4. The advantage of that construction over the much simpler lucky coin-flip construction of Chapter 4, Section 3.2 was that the latter required extreme luck—it required a zero-probability outcome to happen. If we require similar luck to generate our ACCR machine, then that advantage to the more complex construction is lost.

But there is a difference. The lucky coin-flip construction *required* luck: it was only if the array of coins was arranged *by chance* in the lucky way (i.e., with only one heads in each row) that the construction yielded a lottery. In the ACCR machine case, it doesn't matter how the ψ -particle lifetimes or knobs on slices come to satisfy the requisite constraints. Luck is one way in which that could happen. But other options are possible. Maybe we could be in a universe where there is an infinitely complex law of nature that requires the lifetimes or knobs to have the precise values they do, values that in fact satisfy the constraints. Or maybe there is a law of nature that requires the satisfaction of the constraints as such. Or maybe there could be a supernatural being who would choose to make them satisfy the constraints. Or perhaps the Principle of Sufficient Reason is false, and the machine has always existed, for no reason at all.¹⁸

6. Evaluation

A number of mathematical paradoxes follow given AC, and these paradoxes can be operationalized into decision-theoretic paradoxes given Choice Machines as well as, in some of the cases, input from an infinite sequence of coin tosses. It is plausible that if causal infinitism is true, all the constructions involved here are metaphysically possible, and hence we have reason to reject causal infinitism.

¹⁸ There are reasons to think that the PSR supports causal finitism (see Chapter 2, Section 3.2, and Chapter 3, Sections 2.5 and 3.6.2). If so, then the opponent of causal finitism is likely to deny the PSR, and hence it is dialectically acceptable to make use of the denial of the PSR in an argument for causal finitism.

Infinite causal dependencies are involved in two ways with the operationalized paradoxes. First, they are essential to both the functioning and the verification of correctness of a Choice Machine. Second, in the die-guessing paradox and an operationalized coin-flip Dutch Book paradox, we have to be able to take an infinite sequence of coin tosses as an input.

There is another way out of each of these paradoxes. Each paradox relies on some alleged principle of rationality. For instance the version of the guessing paradox that used the Gabay–O'Connor construction in Chapter 4 involved a generalization of the rejection of the Gambler's Fallacy, namely it required the claim that one cannot do better—or at least significantly better—by making use of past information about memoryless dice. The paradox for guessing independent coin tosses assumed a principle about transforming payoffs by permuting equivalent coins. Principles like these could all be challenged. But rejecting causal infinitism gives a more elegant and unified solution to all the paradoxes than rejection of a number of individually plausible principles.

Appendix: **Details of Coin-Toss Rearrangement

To give the details of the rearrangement claimed in Section 5.2 we will use standard methods from Wagon (1994). Let F_2 be the free group on two generators. This is a countably infinite group, and hence there is a bijection ϕ from F_2 onto \mathbb{N} . We can define a group action of F_2 on Ω , the space of countable coin-flip sequences, as follows. For any element g of F_2 , there is a permutation g^\dagger of \mathbb{N} defined by $g^\dagger(n) = \phi(g \cdot \phi^{-1}(n))$ as multiplication by a group member permutes the group. Then let $g\omega = g^{\dagger*}(\omega)$ for $\omega \in \Omega$, so that $(g\omega)_n = \omega_{\phi(g \cdot \phi^{-1}(n))}$.

For any $g \in F_2$, let Ω_g be the set of fixed points under g , i.e., $\Omega_g = \{\omega \in \Omega : g\omega = \omega\}$. I now claim that for any $h \in F_2$ we have $P(h\Omega_g) = 0$ if g is not the identity e . We only need to prove $P(\Omega_g) = 0$, since $P(h\Omega_g) = P(\Omega_g)$, given that the probability measure on Ω is invariant under permutations.

Suppose that $\omega \in \Omega_g$ and $g \neq e$. Then $g\omega = \omega$. Then for all n we have $g^n\omega = \omega$. Write $\omega = (\omega_0, \omega_1, \dots)$. Then $g^n\omega = (\omega_{(g^n)^\dagger(0)}, \omega_{(g^n)^\dagger(1)}, \dots)$. It follows that $\omega_{(g^n)^\dagger(0)} = \omega_0$. Now define $a_n = (g^n)^\dagger(0)$. Then the a_n are all distinct numbers. To see this, suppose that $a_n = a_m$. Then $\phi(g^n\phi^{-1}(0)) = \phi(g^m\phi^{-1}(0))$. Since ϕ is a bijection, it follows that $g^n\phi^{-1}(0) = g^m\phi^{-1}(0)$. Canceling the right multiplicands, we get $g^n = g^m$, and so $g^{n-m} = e$. Thus, either $n = m$ or g has finite order. But the only element g of a free group that can satisfy $g^k = e$ for non-zero k is the identity, so $n = m$. Hence, indeed, the a_n are all distinct.

So, if $\omega \in \Omega_g$, we must have $\omega_{a_0} = \omega_{a_1} = \dots$ for an infinite number of different values of a_n , dependent only on g . But the probability that a particular infinite sequence of coin tosses (defined for a fixed g and h) all comes out the same is zero. So $P(\Omega_g) = 0$. And hence $P(h\Omega_g) = 0$.

Let $D = \bigcup_{g,h \in F_2} h\Omega_g$. Since F_2 is countable, by the countable additivity of classical probabilities, it follows that $P(D) = 0$. And $\Omega - D$ contains no fixed points for any element of F_2 other than e .

By Wagon (1994, Theorem 4.2), there are members α and β of F_2 and disjoint subsets B_1, \dots, B_4 of F_2 whose union is all of F_2 and which satisfy:

$$B_2 = \alpha(B_2 \cup B_3 \cup B_4)$$

and

$$B_4 = \beta(B_1 \cup B_2 \cup B_4).$$

Define the equivalence relation \sim on $\Omega - D$ by saying that $\omega \sim \omega'$ if and only if $\omega = g\omega'$ for some $g \in F_2$. Each equivalence class contains only countably many members of Ω , and Ω has the same cardinality as the reals, so by ACCR, let M be a choice set for the set of equivalence classes under \sim , i.e., a set containing exactly one member of each equivalence class.¹⁹ Then $M \subseteq \Omega - D$ and for each $\omega \in \Omega - D$, there is a unique $\omega' \in M$ such that $\omega \sim \omega'$.

Let $A_i = B_i M = \{g\omega : g \in B_i \text{ \& } \omega \in M\}$.

Observe first that $\Omega - D = A_1 \cup A_2 \cup A_3 \cup A_4$. For fix any $\omega \in \Omega - D$. Then there is an $\omega' \in M$ such that $\omega' \sim \omega$, so there is a $g \in F_2$ such that $\omega = g\omega'$. This g is a member of B_i for some i . But then $g\omega' = \omega$ will be a member of A_i for the same i .

Next, note that the A_i are disjoint. For suppose ω is a member of both A_i and A_j . We need to show that $i = j$. But $\omega = g\omega' = h\omega''$ for some ω' and ω'' in M as well as $g \in B_i$ and $h \in B_j$. Then $\omega' = g^{-1}h\omega''$, so $\omega' \sim \omega''$. Since M contains exactly one element from each equivalence class, we have $\omega' = \omega''$. Thus, $g\omega' = h\omega'$, and so $h^{-1}g\omega' = \omega'$. So ω' is a fixed point of $h^{-1}g$, and since $\omega' \notin D$, we must have $h^{-1}g = e$ and hence $g = h$. But since the sets B_1, \dots, B_4 are disjoint and $g \in B_i$ and $h \in B_j$, it follows that $i = j$. Hence the sets A_1, \dots, A_4 are disjoint.

Now, observe that

$$\alpha(A_2 \cup A_3 \cup A_4) = \alpha(B_2 \cup B_3 \cup B_4)M = B_2 M = A_2$$

and

$$\beta(A_1 \cup A_2 \cup A_4) = \beta(B_1 \cup B_2 \cup B_4)M = B_4 M = A_4.$$

Letting $\rho = (\alpha^{-1})^\dagger$ and $\tau = (\beta^{-1})^\dagger$, we have:

$$\rho^* A_2 = A_2 \cup A_3 \cup A_4$$

and

$$\tau^* A_4 = A_1 \cup A_2 \cup A_4,$$

which is what we were aiming to get.

¹⁹ If f is a choice function, then we can let $M = \{f(A) : A \in S\}$ where S is the set of equivalence classes.

7

Refinement, Alternatives, and Extensions

1. Introduction

We begin by discussing some issues of fine detail needed to refine causal finitism. Are the causes fine- or coarse-grained? Do absences count? With respect to what causal relation are infinite histories ruled out? Among the issues of detail, we will now carefully re-examine the Grim Reaper paradox to see if it is indeed ruled out by causal finitism. No particular refined version of causal finitism will be endorsed, but a number of decision points will be noted for future research.

Next, we will consider alternatives to the causal finitist hypothesis that also kill many or all of the same paradoxes: finitism (already discussed at length in Chapter 1), no regresses, no past infinities, Huemer's no infinite intensive magnitudes, and no room in spacetime. The last of these options offers an answer to an interesting question which we consider on its own: Assuming causal finitism is true, *why* is it true?

Finally, we consider two possible extensions of causal finitism, one to rule out causal loops and the other to rule out infinite explanatory stories.

2. Refinement

2.1 *Event and trope individuation*

Causal finitism says that nothing has an infinite number of items in its causal history. Three main kinds of items are plausible candidates for standing in causal relations: events, substances, and tropes (or particular properties, like the pallor of my face). But both events and tropes present a special problem for causal finitism if they are individuated finely enough.

For suppose that a hot summer day in Texas makes me uncomfortable. If we can individuate events finely enough, we can multiply the causes *ad infinitum*. For instance, suppose it is exactly 38.42° Celsius. Then I am caused to be uncomfortable by its being exactly 38.42°. But at the very same time, I am also caused to be uncomfortable by its being hot, its being more than 38°, its being more than 38.31°, its being around 38.4°, and so on. It seems there are infinitely many causes of my discomfort, contrary to causal finitism.

We are liable to get the same problem with tropes, assuming they have a place in our ontology. The air around me has the particular property of being at 38.42° , and this trope or particular property causes my discomfort. But the air also has the particular property of being hot, of being more than 38° , of being more than 38.31° , of being around 38.4° , and so on, and each of these causes my discomfort.

There are four moves available to get out of this difficulty. First, we can simply accept sparse theories of events and tropes that prevent the multiplication. On the side of events, the usual way to have a sparse theory is to follow Davidson (1963) and take events to be coarsely individuated. Thus the event of its being exactly 38.42° is the same as the event of its being hot and the event of its being around 38.4° , but these events are simply described differently. On the side of tropes, one can deny the existence of tropes like the air's being more than 38° , and insist that the only tropes there are are fundamental ones, perhaps like the air's being exactly at 38.42° , or the air's having such-and-such precise microphysical properties. This move requires no changes to causal finitism, but it will alienate proponents of fine-grained events and abundant tropes.

Second, we can allow for fine-grained events and/or abundant tropes, but suppose that of the multitude of these fine-grained items, only some actually enter into causal relations. If the event of its being exactly 38.42° occurs, then so do the distinct events of its being around 38.4° and its being more than 38.31° , but perhaps out of these only the event of its being exactly 38.42° is causally efficacious, and similarly in the case of tropes.

Third, we could allow that all the multitude of fine-grained items enter into causal relations, but insist that they do so in a way that coheres with causal finitism. For instance, perhaps I am uncomfortable whenever the temperature is at least 34.3° . If so, then it could be that my being uncomfortable is caused by its being at least 34.3° and not by its being exactly 38.42° . On the other hand, there may be a particular level d_{10} of discomfort such that my having exactly level d_{10} is caused by its being exactly 38.42° , and there may be vague levels of discomfort caused by vague temperature ranges. Thus, among the differently granulated causes, only one (or at most a finite number) is causal of an effect of a particular granularity.

These three moves are different ways of denying the infinite multiplication of colocated causes. And there are reasons to deny this multiplication, and hence to opt for one of these moves. One reason is simply that we have good reason to think causal finitism is true, while the infinite multiplication of causes is incompatible with causal finitism. A second reason is the thought that we shouldn't be able to discover an infinitude of physical causes without really serious scientific work—maybe the causal nexus of the physical world *is* infinite in nature, but that fact ought not be obvious. But if we accept the multiplication, an infinitude of physical causes has been discovered just by thinking about a small number of humdrum facts about temperatures and discomfort.

If, however, one likes none of these three moves, there is the fourth option: modify causal finitism. One way to do this is to say that *fundamental* causal histories—causal

histories generated by fundamental causal relations, ones not grounded in other causal relations—must be finite. For instance, even if its being around 38.4° causes my being uncomfortable, it causes it in virtue of its being 38.42° causing my discomfort, and so the former case of causation is not fundamental. Of course, many of the examples of causation in the paradoxes under consideration were likely not to be fundamental either. For instance, in Thomson's Lamp, flipping the switch may only cause the light to turn on or off in virtue of some microphysical instances of causation. However, it is plausible that in all the paradoxes we considered—and unlike in the case of the heat causing my discomfort—the infinite multiplication of non-fundamental instances of causation is paralleled by an infinite multiplication of fundamental instances of causation. Even if none of the switch-flippings are fundamental causes, each switch-flipping's causal contribution is surely grounded in a different collection of fundamental causes.

This gives us a complex decision point for a refinement of causal finitism: *Do we allow for fine-grained events as causes and, if so, do we modify causal finitism to focus on fundamental causes or rule out infinities of fine-grained causes in another way?*

The question of fundamentality makes for a decision point of independent interest: *Do we restrict the causal relations that causal finitism speaks of to fundamental ones?*

2.2 Histories generated by partial causal relations

The rough and ready characterization of causal finitism that I have hitherto used is that causal histories must always be finite. Let's try to make this thesis more precise. It is reasonable to take the causal history of an item z to be the collection of items causally prior to z . But what exactly is this causal priority relation and what are its relata?

Note that to rule out regresses, which is one of the tasks of causal finitism (see Chapter 2), causal priority needs to be a transitive relation: if y is prior to z and x is prior to y , then x had better be prior to z .

A natural option for explicating causal priority is to stipulate that y is causally prior to z if and only if y causes z . The notion of a cause, however, suffers from the infamous selection problem (Hesslow 1988). Normally when asked what caused the forest fire, we would say something like that it was the camper's campfire. But an alien from a planet without much oxygen might equally reasonably say that it was the high oxygen content of our atmosphere that caused the fire. Which counts as the cause seems to depend too much on context.

Moreover, transitivity claims sound dubious if causal priority just is causation. Suppose the camper lit the campfire because she was hungry. Then the hunger caused the campfire. And the hunger was caused by the camper's earlier metabolic processes. But it sounds wrong to say that the camper's earlier metabolic processes caused the forest fire.

The transitivity problem is easily solved. Instead of saying that y is causally prior to z if and only if y causes z , we can say that y is causally prior to z if and only if there

is a chain of causes from y to z , i.e., there is a sequence y_0, \dots, y_n such that y_i causes y_{i+1} for each i and $y = y_0$ while $z = y_n$.

Even so, partial rather than full causation is enough for some of our paradoxes (and the others can perhaps be modified to use partial causation). For instance, take the paradox from Chapter 5, Section 2.5 where an angel announces that among the infinitely many dice, only finitely many showed something other than six. To apply causal finitism, I noted that the infinitely many states of the dice caused the angel's announcement. But no one state of the dice was a full cause of the angel's announcement. In fact, no one state of the dice made any difference to the angel's announcement: if the die had come up otherwise, with all the other dice fixed, there would still have been only finitely many non-sixes. It seems that there is a full cause, and that's a plural cause: the totality of all the die rolls. But that plural cause is only one, and one is a finite number, so if causal finitism applies only to full causes, it does not apply here.

Perhaps this is too quick. Maybe there are infinitely many full causes: the state of all the dice but the first, the state of all the dice but the second, the state of all the dice but the third, etc. But we should not rely on that. Maybe what we have is simply a direct interaction of all the dice taken together with the angel's mind. And we need causal finitism to rule out such cases, too.

It is better, thus, to work with a weaker concept than causation to define causal priority. Partial causation or causal contribution is such a notion. We could thus say that x is causally prior to y provided that there is a chain of partial causes from x to y . I shall use the terms "partial causation" and "causal contribution" interchangeably, though future research might find subtle differences to distinguish them.

Thus, if we follow this suggestion, causal finitism will be the thesis that there are no infinite histories generated by causal priority, defined in terms of chains of partial causation.

This leads to a second decision point: *Does causal finitism apply only to full causation or also to partial causation or causal contribution?*

Here, I think the answer should be affirmative. Furthermore, fine analysis of the Grim Reaper paradox may suggest that we need an even weaker relation than partial causation.

2.3 A closer look at Grim Reapers

Recall the Grim Reaper story from Chapter 3, Section 3. The lamp is off at 10 am. At an infinite sequence of times t_1, t_2, \dots strictly between 10 and 11 am the Grim Reapers activate. When a Reaper activates, it checks if the light is on. If it's on, it does nothing. If it's off, it flips the switch to turn the lamp on. I have argued that the story becomes paradoxical when $t_1 > t_2 > \dots$ (e.g., if the n th reaper activates $30/n$ minutes after 10 am), but that we should rule out the story no matter how the t_n are arranged on the grounds that it is a violation of causal finitism.

But now consider a setup where the activation times are ordered in the opposite way, so that $t_1 < t_2 < \dots$ (e.g., the n th reaper activates $30/n$ minutes before 11 am). Then there would be no paradox: the Reaper that activates at t_1 turns the lamp on, and the subsequent ones have nothing to do. Where is there an infinite causal history here? Only one Reaper actually affects the state of the lamp in this story. And in the paradoxical version where $t_1 > t_2 > \dots$, no Reaper turns on the lamp. Where is there infinite causality here? It appears, then, that causal finitism does not rule out the Grim Reaper scenario and its cousins as claimed, and hence the philosophical utility of causal finitism is significantly reduced.

There are at least six solutions to this, each leading to a different refinement of causal finitism. The first solution is to note that what we have is an infinity of items that each have the *power* to light the lamp. We could extend causal finitism to rule out not only actual infinite cooperation but also cases where an infinite number of things each has the power to contribute to some effect (where the effect is identified in a way that doesn't require essentiality of origins, perhaps by qualitative features and spatiotemporal location), even if they do not actually do so.

This is an attractive option to consider as it rules out some further cousins of the Grim Reaper story. For instance, take the Grim Reaper paradox but simply remove the lamp, and specify that if a Reaper sees no lamp, it does nothing; thus, each Reaper wakes up, does nothing, and goes back to sleep. It would be odd if one could have the infinite number of Reapers arranged as in the paradox without the lamp, but it would be impossible to add the lamp. The power approach has the advantage of allowing us to directly rule out this story, as we have an infinity of items that have the power to light any lamp that's there. Without bringing in powers, we may have to instead say that the story is ruled out because if it were possible, it would be possible to add a lamp, and that would be ruled out by causal finitism.

The second solution is to extend the scope of the partial causation relation to include what *each* Grim Reaper does for the lamp's final state. In an international survey of jokes, the following morbid joke by Gurpal Gosall was rated as the funniest:

Two hunters are out in the woods when one of them collapses. He doesn't seem to be breathing and his eyes are glazed. The other guy whips out his phone and calls the emergency services. He gasps, "My friend is dead! What can I do?" The operator says "Calm down. I can help. First, let's make sure he's dead." There is a silence, then a shot is heard. Back on the phone, the guy says "OK, now what?" (Wiseman 2002)

The shot may or may not have caused death, but it *ensured* it. In moral action, ensuring has many of the same normative consequences as doing. But by his shot, the hunter made himself into the moral equivalent of a murderer, whether or not the companion was already dead. Note, too, that shooting the companion wasn't the only way to ensure he was dead. The other way—probably legally safer but still morally depraved—would be to adopt the plan of checking whether his companion was dead and shooting him if and only if he wasn't.

Now, the verb “to ensure” suggests agency. But we can use it in non-agential cases as well: “The avalanche at 11 am ensured that Smith was dead by noon, but to figure out the disposition of his estate it was important to figure out whether Smith was dead before 11 am.” We can recognize ensuring as a kind of causal relationship which can be agential as well as not.

Given this background, we could understand the partial causation or causal contribution relation that generates the causal histories that according to causal finitism are finite to include not just cases of what one might call actual causation but also cases of partial or contributory ensuring. And the Grim Reapers all engage in that, even if they do nothing but check if the lamp is on. They are like the legally more cautious murderous hunter who checks if his companion is dead and shoots only if he isn’t. We can even understand such a partial or contributory ensuring as being indeterministic in nature—imagine that the murderous hunter shoots his companion in a way that does not guarantee death but happens to succeed.

There is something awkward about making the causal history of the lamp’s being on at 11 am include Grim Reapers that didn’t touch the switch. If so, then perhaps the term “causal history” is not the best one to choose. Perhaps one could say “backwards causal nexus” instead.

The third solution is to extend causal finitism to forbid not just causal dependence but also causally-grounded counterfactual dependence on an infinite number of distinct positive items (substances or events) such that had none of these events happened, the caused event wouldn’t have happened. We distinguish this from logically (or metaphysically) grounded counterfactual dependence: my being six feet tall is counterfactually dependent on my being more than three feet tall (if I weren’t at least three feet tall, I wouldn’t be six feet tall), and on an infinite number of other such properties, but the counterfactual is grounded logically rather than causally. And we also distinguish causally-grounded counterfactual dependence from counterfactual dependence on negative items—that I am writing this is counterfactually dependent on an infinite number of absences of witches that could have prevented my writing. And then we say that the lamp’s being on is counterfactually dependent on an infinite number of Reapers, and hence the story is ruled out by causal finitism.

Perhaps, in fact, causally-grounded counterfactual dependence is a kind of causation—if so, then we have not really extended causal finitism but merely clarified it.

The fourth solution is inspired by some remarks of Philip Swenson. We could say that an item counts as causally impinging on an event provided that it actually causally contributed to (or partially caused) the event or would have causally contributed to (or partially caused) whether the event would have happened but for something else that happened. Each Grim Reaper would have causally contributed to the lamp’s being on at the end of the experiment, but for the previous ones. And now our causal finitism needs to rule out infinite histories of causal impingers.

The fifth solution is to take there to be an intuitive notion of a transitive relation of causal priority that is either not further analyzed or is not further analyzable. It is intuitively compelling that the Grim Reaper paradox, and our other paradoxes, are ruled out if we forbid items from having infinitely many items causally prior to them. But as the above discussion shows, it is difficult to spell out what the causal priority consists in. So, the suggestion is that we not spell it out!

The sixth solution is this. The final state of the universe at 11 am is causally affected by each of the infinitely many acts of observation of the lamp. Each Grim Reaper observes the state of the lamp, and consequently does or does not push the button on the lamp. Not pushing the button certainly affects the final state of the universe. Had a Reaper pressed the button, the universe would have been in a different state at 11 am. The button might have, for instance, had a Reaper fingerprint (or bone print!) on it; the change in the position of the Reaper's arm would mean that the gravitational field spreading at the speed of light would be forever different; and so on.¹ We may not want causal finitism to *simply* rule out infinitely many absences that are causing by omission (we will discuss absences in Section 2.5). But a positive cause can act *by means* of an absence. For instance, one method of killing an animal is by depriving it of air—yet that is clearly and literally an instance of something positive, say, strangulation, causing something else, namely death. Cases where infinitely many positive events—in this case, observations—affect things *by means* of absences (say, of button pressing) can be taken to be ruled out by a literal and unextended reading of causal finitism. The Grim Reaper story is like that: the final state of the universe is causally affected by each of the observations—had any observation come out differently, the final state of the universe, not in respect of whether the light is on or off but in respect of subtler things, would have been different.

This sixth solution is an attractive one, but it may not handle modifications to the paradox. We could imagine a world where Reapers leave no fingerprints on the button, where gravitational fields don't spread from their arms' movements, and so on. Or a world where everything is destroyed at 11 am (even the switch) except the shining or non-shining lamp. I am not sure if this objection to the solution works, but my uncertainty depends on a controversial view of time. Following Kant (1907) (see also Carrier 2003), one may think that the direction of time is constituted by causal relations. On this view, it may well follow that if a Reaper's act of observation is not in any way causally prior to the state of the universe at 11 am, then its temporal location cannot be said to be prior to 11 am either.

To illustrate the six solutions, let us apply them to Benardete's Boards paradox (Benardete 1964, pp. 237–8), which can be thought of as a non-temporal variant of the Grim Reaper paradox. We will find that the first five work in that case, but the last one has more difficulty. There are infinitely many impermeable boards perpendicular

¹ Cf. Lewis (1979).

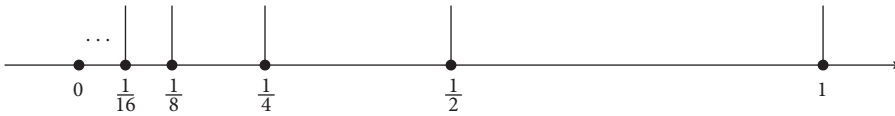


Fig. 7.1 Benardete's Boards.

to the x -axis, say at coordinates $1, 1/2, 1/4$, and so on. The boards are either infinitely thin or else their sizes are such as to allow a space between them (e.g., the board at 2^{-n} has thickness 2^{-n-2}). A point particle moves along the negative x -axis in the positive direction. That particle cannot go past $x = 0$ because to go any distance beyond $x = 0$, it would have to cross infinitely many boards. But which board stopped the particle? It seems that none did: the board at $x = 1/2$ did not, since the particle never reached it; neither did the board at $x = 1/4$, and so on (Fig. 7.1).

But we can say that an infinite number of boards each had the power to stop the particle from reaching, say, $x = 1/2$. Likewise, we can say that an infinite number of boards *ensured* that the particle did not reach that coordinate. We have counterfactual dependence on the stopping powers of infinitely many boards: had none of the boards been there (or had they been permeable), the particle wouldn't have stopped. And each board would have contributed to the particle's stoppage but for the others. Moreover, the causal priority of infinitely many items is just as plausible in this case as in the Grim Reaper case. Hence the setup is ruled out by the five improved versions of causal finitism.

However, the sixth solution, which did not require any rewording of causal finitism as such, does not appear to generalize to Benardete's Boards. There does not appear to be any analogue to the acts of observation in the case of Benardete's Boards. So if we opt for that sixth solution, we have to admit that there is a causal paradox that causal finitism as such does not directly rule out. Though it might indirectly rule it out, if causal finitism forces space to be discrete, a question we will explore in Chapter 8.

This yields a third decision point: *Do we need to weaken the causal relation further so as to make causal finitism kill the Grim Reaper and Benardete's Boards paradoxes, and if so, how?*

2.4 Objections to causal finitism involving partial causation

But now we need to be careful that we don't make causal finitism get too close to finitism. Suppose we take partial causation as the relevant relation. Then, plausibly:

- (1) If x causes y and y is a part of z , then x is a partial cause of z .

But now consider an unparadoxical forwards causal sequence: x_1 causes x_2 , x_2 causes x_3 , and so on. Let z be the fusion of all the x_i . Then each x_i is a partial cause of z by being the cause of x_{i+1} which is a part of z . So z has infinitely many partial causes. But we shouldn't rule out unparadoxical forwards causal sequences. The same argument can be run in terms of causal contribution.

This doesn't necessarily refute the claim that the relevant causal relation for causal finitism is partial causation. Perhaps we should reject (1) instead. After all, it is plausible that such a forwards causal sequence is possible. It is even plausible that there is a world where such a forwards causal sequence constitutes all of concrete reality and where x_1, x_2, \dots are simples. But it is not plausible that a fusion of simples could be even partially caused by each of the simples composing it, as that would be too much like self-causation. Yet z , the fusion of the x_i , would precisely be a fusion of simples partially caused by all the simple parts if (1) were true.

Alternately, we could endorse (1) but deny that there are such infinite fusions as the fusion of all the x_i , either because of a general skepticism about fusions or because of a skepticism about infinite fusions. One might think of the parts of a fusion as Aristotelian "material causes". Then denying infinite fusions would be a natural extension of causal finitism (see Section 5.2).

Here is a related worry. Suppose there is an infinite number of causally disconnected universes, and suppose that in the n th universe there is an object x_n which causes an object y_n . Consider the fusion Y of all the y_n . Then the causal history of Y includes infinitely many objects: x_1, x_2, \dots . Should we take causal finitism to rule out the above scenario? If causal finitism rules out an infinite number of causally disconnected universes, then we are well on our way to finitism as such.

There are, however, at least two ways of maintaining causal finitism without denying the possibility of an infinite multiverse. The first is, once again, to turn the above remarks into an argument against unrestricted fusions. Perhaps wholes have to have some sort of organicity or at least causal interconnection, or perhaps infinite wholes are impossible.

The second way is to deny that x_n is an element of the causal history of Y . While there are cases where the cause of a part is a partial cause—any German officer who ordered an initial part of the invasion of Poland was a partial cause of World War II—in general a cause of a part need not be a partial cause. The D-Day invasion was not a partial cause of World War II, even though it was a cause of many battles that are a part of World War II. Nor would a mad scientist who made my body sprout a tail be a partial cause of me, even though she would be a cause of a part of me. Thus, perhaps we should instead say that the cause of Y is, rather, the fusion X of all the x_n , and the individual x_n are not elements of the causal history or nexus, even though they are causes of parts of Y .

We thus have a choice point: *Do we rule out infinite fusions in addition to infinite causal chains or do we carefully analyze the relevant causal relation in order to rule out counterexamples coming from such fusions?*

2.5 Absences and omissions

Someone who suffocates dies from the absence of oxygen. Should we include such absences in causal histories?

The natural answer is that if absences are causes, then they had better be included in causal histories. But the worry about this is that just about every event then has infinitely many causes. A causal contributor to my writing this book is the absence of an earth-destroying attack by bipedal aliens ten years ago. Another contributor is the absence of such an attack by tripedal aliens, and another is the absence of such an attack by quadrupedal ones. There are infinitely many absences of genocidal alien invasions that have contributed to this book, and so causal finitism has no hope of being true if we include absences.

One solution is simply to stipulate that it is only positive items that count for causal finitism.

The other solution is to generate causal histories by means of fundamental cases of the relevant causal relation, as has already been suggested in Section 2.1. And it is implausible that the causal relationship of each of the infinitely many items of the form *the absence of an earth-destroying attack by n-pedal aliens ten years ago* to my writing of this book is fundamental. Perhaps what is fundamental is the contribution of some more general absence, such as the absence of any earth-destroying attack prior to the present, or maybe the absence of any event that significantly impedes my normal functioning, or the like. Or perhaps few or no absences are involved in any fundamental case of the relevant causal relation.

We have, thus, a choice point: *Does causal finitism apply to causation by absences?*

Note that even if causal finitism doesn't rule out causation by an infinite number of absences, it probably should rule out an infinite number of positive events contributing to an event *by causing an intermediate omission or absence*, as we argued in one of the solutions in Section 2.3. There is a difference between an absence being a cause and a positive event causing by means of omission. When a guillotine's fall causes someone's death, that's a case of a positive event causing death by means of the absence of oxygen in the brain. An infinite number of guillotines causing a single death should be ruled out by causal finitism.

One may worry, too, that death itself is an absence of life. There will also be a choice point here: *Does causal finitism apply to causation of an absence?* If not, then we might rule out an infinite number of guillotines causing a single death by noting that in that case the infinite number of guillotines would also cause a movement of the head, or consequent decay, or the like. But it is neater to apply causal finitism directly to the causation of an absence.

Another possibility is that even if we do not see causal finitism as forbidding causation by an infinite number of mere absences, it could forbid causation by an infinite number of *privations*. Haldane (2007) has argued that even if mere lacks can't be causes, privations can be. It is incorrect to say that Smith's lack of wings caused him to be late for a meeting, even if he would have been on time had he wings, but it could be correct that Smith's being deprived of legs caused him to be late for the meeting. This might mesh nicely with one of the solutions to the Grim Reaper problem. For it

might be that when something has a power (or at least disposition) to do something on a triggering condition—as when each Grim Reaper has the power (and disposition) to turn on the lamp on the condition that the lamp is off—then it counts as deprived of the exercise of that power when the condition is absent.

3. Some Competitors to Causal Finitism

3.1 *Finitism*

It is time to consider other hypotheses that rule out paradoxes of infinity. The first of these was already considered in Chapter 1: finitism. Finitism is simpler than causal finitism and that makes it an attractive competitor. Moreover, it has the advantage of ruling out non-causal paradoxes of infinity. However, as we saw it also rules out too much: it is difficult to reconcile with mathematical facts such as the existence of infinitely many primes—or other numbers.

Furthermore, the simplicity of finitism is somewhat deceptive. For it is only when married to further metaphysical theses that finitism rules out the paradoxes we have considered. If presentism is true, then finitism only tells us that there is no *present* infinity of items, and that's not enough to rule out many of the causal paradoxes of this book. To rule those out, finitism needs either eternalism or a growing block theory of time.

Given eternalism, past, present, and future infinities are ruled out by finitism. But the resulting theory is very counterintuitive: it rules out the possibility of an infinite future collection of events. This means that finitism, in order to both do its paradox-busting work and avoid refutation by reference to future events, must be attached to a growing block theory of time.

Now, while finitism by itself is a simple theory, finitism plus growing block is more complex, and inherits the well-known difficulties of growing block. For instance, first, growing block theories posit an absolute simultaneity: all events at the leading edge of the growing block are absolutely simultaneous. While with some work this can be reconciled with the empirical predictions of Relativity Theory (Smith 1993, Section 7.2), it is still in tension with the metaphysics that most naturally arises out of relativity, a metaphysics on which there is no metaphysically absolute simultaneity.

Second, in Chapter 1, Section 4.1, we have discussed the argument of Merricks (2006) that growing block leads to an absurd skepticism about whether this time is present.

Finitism is not, thus, a good competitor to causal finitism. It is in tension with mathematics, and we would need to add to it the growing block theory of time which increases the overall complexity and vulnerability of the theory. Let us now consider some further competitors.

3.2 *No infinite regresses*

As seen in Chapter 2, there is intuitive support for denying the existence of infinite causal regresses independently of the various paradoxes. Further, denying the existence of infinite causal regresses will kill a number of paradoxes. For instance, the natural way to understand the die-guessing paradoxes of Chapter 5 involves an infinite regress. At time 0, your knowledge of the rolls at times $\dots, -4, -3, -2, -1$ is partly caused by your knowledge at time -1 of the rolls at times $\dots, -4, -3, -2$, which in turn is partly caused by your knowledge at time -2 of the rolls at times $\dots, -4, -3$, and so on.

However, supertasks of the Thomson's Lamp variety do not actually involve an infinite regress: prior to each flipping of the switch there are only finitely many flippings. Further, there are paradoxes involving an infinite number of causes working together but apparently not in a regressive fashion, as in the interpersonal die-guessing game in Chapter 5, Section 2.5. So while a no-regress thesis will do a lot of the same work as causal finitism, it does not do enough of the work.

3.3 *No past infinities*

There is an Aristotelian intuition whose slogan is that *completed* infinities are impossible (e.g., see Craig 2009). One cannot *traverse* an infinity. The infinity of future days is certainly possible, but even if there is an infinite afterlife, no one will ever be able to say truthfully: "I have completed an infinite life." Thus the intuition does not rule out a future infinite.

It is natural to read the Aristotelian intuition as a denial of the possibility of *past* infinities. A past infinity is completed, done with. The world has traversed it. And that's impossible according to the Aristotelian.

If causes must be temporally prior to their effects, then the impossibility of past infinities immediately entails causal finitism: if there is an infinite causal history, then at the time of the effect there is an infinite past.

But it is plausible that causes can also be simultaneous with their effects, as in Kant's (1907) famous example of an iron ball simultaneously depressing a cushion.

If simultaneous causation is possible, then the above argument for causal finitism from the impossibility of past infinities fails in the case of an infinite number of items simultaneously causing something, when all this causation is happening at the very last moment of time. However, it is extremely plausible that if an infinite number of items can simultaneously cause something at the last moment of time, they can simultaneously cause something at a moment that is succeeded by a later moment. So we still have a strong argument from the impossibility of past infinities to causal finitism if we weaken the assumption that causes are earlier than their effects to the assumption that causes are earlier than or simultaneous with their effects.

Unfortunately, even the weaker assumption may be false. First, perhaps something temporal could have an atemporal cause. Classical theism holds that God is atemporal and yet he created a world with time. Second, apart from cases where backwards causation gives rise to causal circles, backwards causation—where the cause is later than the effect—appears to be metaphysically possible, even if it may be contingently ruled out by the laws of physics.

However, if causes don't need to be earlier than or simultaneous with their effects, then the denial of past infinities is insufficient to rule out all the paradoxes we need to rule out. For instance, consider the guessing game involving infinitely many people in Chapter 5, Section 2.5. What was crucial there was the availability of information about infinitely many die rolls. How the rolls are arranged in time is irrelevant. The story would work just as well, for instance, if each roll happened in a separate island universe, without there being a single time series running through all the island universes.

In fact, if causes do not have to be earlier than or simultaneous with their effects, then very plausibly time travel is possible. But if time travel is possible, then it should be possible to live one's life in reverse, being conceived and born in the (external) future and living to the past. The easiest way to get that possibility is to suppose time is a series of discrete moments. Our backwards-living person then time-travels from each moment to the previous, instead of "time-traveling" from each moment to the next as we all normally do. But if such a backwards life is possible, then the denial of past infinities is insufficient to rule out the possibility of traversing an infinite series. For imagine someone who, according to our external time, now exists and will always exist. But she lives her life backwards. This person, thus, at this point can be said to have lived an infinite life. And if this is her moment of death, then she has completed that infinite life, contrary to the Aristotelian intuition that there are no completed infinities.

So if causes do not have to be earlier than or simultaneous with their effects, the very Aristotelian intuition that pushed one to deny past infinities also pushes one beyond that denial. Indeed, it pushes one to causal finitism, because as we saw in the backwards-life case, what is problematic for the intuitions is the existence of infinitely many causes, not their temporal arrangement.

But once we have got here, and admitted that we should accept causal finitism, it is worth asking whether we need to deny past infinities as such. Causal finitism appears sufficient to do justice to many of the intuitions about traversal of an infinite sequence. And it is the admixture of causation that seems to make for real paradoxes.

We can, thus, take causal finitism to be the best way to do justice to the intuitions behind the denial of completed infinities. The completed here can be understood as what is ready to hand (to put it in Heideggerian terms!), what is in some sense usable, namely what is ready to be a cause, and in this sense, if causal finitism is true, there can be no completed infinities.

3.4 No infinite intensive magnitudes

3.4.1 THE BASIC THEORY

In a recent book, Huemer (2016) has surveyed a collection of paradoxes of infinity overlapping the collection considered in this book, and employed the same strategy of finding a metaphysical hypothesis that rules out paradoxes. In his case, the hypothesis was that there can be no infinite intensive magnitudes.

The magnitudes Huemer is interested in are all natural magnitudes, the kinds of magnitudes that can be found in scientific explanations (cf. Huemer 2016, pp. 135–7). For instance, if m is the mass of an object in grams and h is the volume of the object in cubic inches, then $m + 2^h$ is not among the natural magnitudes Huemer is interested in. On the other hand, magnitudes like charge, mass, and volume are natural magnitudes.

Now, these natural magnitudes divide into two kinds. An extensive magnitude is one that can be defined by adding up values of some natural magnitude. For instance, the total mass or charge of an object is the sum of the masses or charges of the parts. Magnitudes that are not defined as such sums are intensive. Huemer allows some extensive magnitudes to be infinite, but intensive ones can never be. Thus, it is in principle possible for the universe to contain an infinite amount of mass, as the total mass of the universe is an extensive magnitude. This is a good thing: it fits with the intuition that the universe might be infinite, and allows Huemer to reject finitism, just as this book does. On the other hand, it is not possible for the universe to have infinite density according to Huemer, since density is an intensive magnitude, being the *ratio* of two sums (the total mass and the total volume) rather than a sum.

The restriction of the theory to natural magnitudes is necessary. For if we allow non-natural magnitudes, then finitism immediately follows. For if finitism is false, then the log-count of objects in existence is infinite, where the log-count of F s is the logarithm of the number of F s, but is not defined as a sum. Fortunately, the log-count of objects in existence doesn't enter into scientific explanations.²

This allows Huemer to rule out a number of paradoxes by arguing that plausible embodiments of these paradoxes will involve an infinite intensive magnitude. For instance, the total distance moved by the toggle switch in Thomson's Lamp will be infinite. That's fine, since total distance moved is an extensive magnitude. But the speed of toggle switch motion will also be infinite, since it will be given by D/T where D is the infinite total distance and T is the finite amount of time that the toggles took.

² I assume that the number of F s can be considered a sum and hence is extensive: it is the sum of one unit per F . This is a little awkward, but Huemer has to say something like it since the number of objects of some natural kind does enter into scientific explanations—for instance, the number of organisms in a kind is studied by population biology—and it would be uncomfortably close to finitism to have to say that no natural kind can have infinitely many things.

And while both D and T are extensive, D/T is not, so Thomson's Lamp involves an infinite intensive magnitude, and is thus impossible.

Huemer's resolution of the paradoxes has multiple problems. First, I will argue that by Huemer's account of intensiveness, there are several examples of intensive natural magnitudes that can be infinite, assuming that finitism is false. Second, Huemer's strictures against intensive infinities do not in fact kill all the paradoxes he applies them to, because some of the intensive infinities in question are not natural. Third, I will discuss the possibility of immaterial beings that slip through Huemer's strictures.

3.4.2 SOME INFINITE INTENSIVE MAGNITUDES

3.4.2.1 Center of mass and moments of inertia The center of mass of a plurality of objects is a natural (vector-valued) magnitude. But if finitism is false, it is surely possible to have a stack of pancakes of equal size and mass, with the stack being infinite in the upward direction—let's suppose that's the z -axis—and finite in all the other directions. The z -coordinate of the center of mass of the stack of pancakes will then be infinity. But centers of mass are not extensive: they are weighted averages of positions rather than sums. Hence, if finitism is false, it is possible to have an infinite natural magnitude that is not extensive.

One might object that Huemerian magnitudes must be scalars rather than vectors. If so, then simply run my argument in one dimension. One could have a universe with only one spatial dimension, and the center of mass in such a universe could be a scalar (measured relative to a privileged zero point, say). And we can simply imagine a sequence of massive particles going off to infinity in one direction but not the other, which would result in such an infinite position of the center of mass.

3.4.2.2 Mental life Jim doesn't like having hair, and assigns a utility to having n hairs that is proportional to $1/n$, with 0 hairs getting infinite utility. Jim ends up plucking out all of his hairs as a result of his utility assignment. The explanation for why Jim plucked out all his hairs is that his subjective utility for that state of affairs was infinite. But the subjective utility for a state of affairs is a natural magnitude: it enters into explanations in economics and psychology. Moreover, the subjective utility for a state of affairs is not, in general, an extensive magnitude. Hence, there is an infinite natural intensive magnitude, contrary to Huemer.

Human beings are the object of scientific study and yet are capable of making all sorts of magnitudes be explanatorily relevant, and hence natural in the sense needed for Huemer's account. Perhaps, though, Huemer might argue that psychology is not a fundamental science, and it is only *fundamental* non-extensive infinite magnitudes that are impossible. But this undercuts Huemer's own applications of his theory to paradoxes. Overall speeds, densities, and strengths of materials are not a part of *fundamental* physics. They are at best parts of a higher-level approximately Newtonian physics. Yet Huemer needs to count such magnitudes as intensive and ruled out by his theory, despite their non-fundamentality.

3.4.2.3 *Black holes* Huemer (2016, p. 159) rejects black holes because of the infinities involved. Here we presumably have to distinguish between black holes as such and black holes as described by Relativity Theory. It is well-established astronomy that there are black holes, for instance a supermassive one at the center of the Milky Way galaxy located at the radio source Sagittarius A*. However, whether these objects are correctly described by Relativity Theory is another question. If Huemer is not to dispute the observational evidence, he has to hold that intensive infinities that one gets in the relativistic description of the object at Sagittarius A* objects are a mark of the incorrectness of the description, and indeed of the metaphysical impossibility of that description being correct. That is a costly move.

3.4.2.4 *Particles* Density is a natural magnitude for Huemer (see, e.g., Huemer 2016, p. 211). This immediately rules out point particles that have non-zero mass. Electrons have mass, so Huemer has to think that either they have no meaningful dimensions or else their size is non-zero.³ This puts a significant constraint on particle-based physics. It means that fundamental physics based on point particles with mass is not only false, but could not possibly be true.

Interestingly, Huemer may also have to reject particles with mass that have a non-zero size. The reason for considering densities to be natural magnitudes is that they enter into explanations in (non-fundamental) science. Likewise, density gradients enter into scientific explanations. Here is a description of “Density Gradient Centrifugation”:

We can make use of the dependence of sedimentation velocity on solvent and particle density [...] by centrifuging particles through a medium of *gradually increasing* density. This is achieved by establishing a *density gradient* between a region of high density at the bottom of the centrifugation tube and a region of low density at the top [...].

(Sheehan 2009, Section 7.5.2, emphases in original)

A density gradient is the rate of change of density as one moves through space. It is clearly not an extensive magnitude.

But now, plausibly, if there can be particles of non-zero size and non-zero mass, there can be particles of non-zero size, well-defined spatial extension, and non-zero *uniform* density ρ . If we put such a particle in vacuum and consider a path from vacuum into the particle, along the path we would find an instant change in density from zero to ρ . This would result in an infinite density gradient along the path.

Thus, the denial of intensive infinite magnitudes will run into problems with both point *and* extended particles of non-zero mass. It will likewise have problems with any discontinuous transitions in density, say between two uniform non-particulate materials of different density. And what goes for mass density and mass density gradient will apply to charge density and charge density gradients: thus, there will be problems with point and extended particles of non-zero charge.

³ I am grateful to Ian Slorach for impressing on me the importance of this.

Huemer's best bet would probably be to insist that fundamental particles cannot possibly have a well-defined spatial extension. On some interpretations of Quantum Mechanics that might be true, but this is one of the *counterintuitive* consequences of Quantum Mechanics. It is a cost for Huemer not only to have to insist that this counterintuitive consequence is true, but that it is necessarily true.

3.4.3 HUEMER'S INTENSIVE MAGNITUDES

3.4.3.1 *Speed, Thomson's Lamp, and Hilbert's Hotel* The toggle switch in Thomson's Lamp moves an infinite distance. This is extensive. But overall speed is total distance traveled divided by time, and hence is intensive. Hence, Thomson's Lamp involves an intensive infinite.

However, Huemer does not rule out all intensive infinities, but only the natural ones, the ones that enter into explanations. Instantaneous velocity definitely enters into scientific explanations. Maybe instantaneous speed—the magnitude of the velocity vector—does as well. But it is far from clear that overall or average speed⁴ is a natural magnitude. It is tempting to say that the fact that an object moved over a distance D in a period of time T is explained by its having an average speed of s such that $D = sT$. But it seems more correct to say that the explanation for why the object moved D in T is that D equals the integral of the instantaneous speed over the period of time. After all, we shouldn't say that an elevator has 300 kilograms in it because it contains four people whose average mass is 75 kg, but rather that the elevator has 300 kilograms in it because that is the sum of the masses of the people in it—there is no need to divide the total by 4 to get the average and then multiply by 4 again to get the explanation.

Moreover, there is a reason why specifically Huemer should not count overall speed as an intensive natural magnitude. Imagine a particle in a non-relativistic world that accelerates exponentially so that in the n th second it travels 2^n meters. This particle's average or overall speed over all future time will be infinite. Huemer (2016, p. 160) suggests that there may have to be a maximal limit to the speed of objects according to the laws of nature, but this simply is not plausible. Doing so would make Newtonian mechanics metaphysically impossible, because if we had a Newtonian disc whose circumference moved at the maximum speed, we could stand on the disc, extend a stick past the periphery, and the tip of the stick would move at a higher speed (cf. Leibniz 1989, p. 237).

Additionally, it certainly seems metaphysically possible to have an immaterial mind or a magician that can instantaneously teleport a switch from one position to another.

⁴ **The two are the same assuming the movement is sufficiently smooth. If an object moves from time 0 to time T and its position vector at time t is $x(t)$, then assuming the position is differentiable and the derivative is integrable, the distance traveled will be $\int_0^T |x'(t)| dt$, so the overall speed will be $(1/T) \int_0^T |x'(t)| dt$. But the average speed will be the same, since $|x'(t)|$ is the speed at time t .

Huemer has some other moves here. The amount of work expended in flipping the switch is infinite, which he claims will generate a black hole (Huemer 2016, p. 198). Huemer, however, believes that black holes as understood in General Relativity are impossible as they involve infinite densities and curvatures, and concludes from this that General Relativity is false (Huemer 2016, p. 159). This creates a dilemma for Huemer. If General Relativity is true, Huemer's theory fails by his own admission. But if General Relativity is false, what guarantee is there that an infinite amount of work would produce a black hole? And, in any case, the claim that an infinite amount of work would have to be involved is dubious. We could suppose, for instance, that the mass of the switch and the friction of its bearings decreases with each flip (this also takes care of the worry in Huemer 2016, p. 198 about infinite heat being produced from the system).

3.4.3.2 *Smullyan's rod Recall how causal finitism ruled out Smullyan's rod, the semi-infinite rigid rod balanced on one end some distance above an infinite planar surface because it cannot swivel down by any angle, no matter how small (Chapter 3, Section 4.2). Huemer (2016, p. 184) has a different resolution. There will be infinite forces involved. The rod being infinitely heavy will have to bend unless it has infinite resistance to bending. It will, for the same reason, puncture the plane it is suspended above unless that plane has infinite strength. But strengths are intensive and hence cannot be infinite.

However, the story can easily be modified to avoid infinite forces. Suppose that density in each meter of the rod, starting at the finite end, is half of the density in the preceding meter. Then the total mass of the rod is finite, the gravitational forces are finite, and the gravitational torque about the pivot or any point in the rod will also be finite.

But even with finite forces there is the problem of the rod's rigidity. Let's consider this problem with a little more detail than Huemer does. In the real world, a solid isotropic rod of circular cross-section of diameter D and finite length L fixed at one end with a perpendicular load force F at the other end exhibits (at least approximately and for forces small enough to be in the linear regime) a deflection of $\frac{64FL^3}{3\pi D^4E}$ (see Nielsen and Landel 1994, p. 38). Thus, as long as the force is non-zero and the Young's modulus is finite, the deflection is non-zero. But any deflection will bend the rod, and the result will presumably be that some distance away from the pivot the rod will be lying along the ground (option (d) in Huemer 2016, p. 184).

Thus, while Huemer's problem of infinite forces can be avoided, the rod would need to have infinite Young's modulus to avoid bending even under finite forces. Huemer can now say that Young's modulus is an intensive quantity and hence cannot be infinite. I could, of course, grant that Huemer's solution works for Smullyan's rod. But things are not so clear.

First, it is not plausible that the kind of proportionality of deflection to small forces that the $\frac{64FL^3}{3\pi D^4E}$ formula embodies is metaphysically necessary. It seems metaphysically

possible to have a classical physics world where instead of linear bending responses to small forces, there is a threshold response so that for small enough forces, there is no bending at all, and bending only begins once the force reaches some threshold. For instance, the deflection law could say that the deflection is 0 if $F \leq F_0$ where F_0 is determined in some nomic way from the dimensions of the rod and is $\frac{64(F-F_0)L^3}{3\pi D^4 E}$ for $F \geq F_0$.

Thus small tweaks to the paradox and the laws make Huemerian resolutions fail. But the causal finitist resolution, where the magnitude of torque and degrees of rigidity were irrelevant, is unaffected.

Second, it is not clear why Young's modulus rather than its reciprocal is in fact the natural magnitude. We can think of Young's modulus for a material as determined by the limiting ratio of compressive force to deformation (cf. Nielsen and Landel 1994, p. 36):

$$E = \lim_{L \rightarrow L_0+} \frac{F/A}{(L - L_0)/L_0},$$

where we have an upright cylinder of material with initial height L_0 and a top face of area A and a force F is applied perpendicularly downward to the top face of the cylinder to squash it to height L . Thus, E is a measure of the strength of the material. But one could have equally well defined something like the propensity E' to compress:

$$E' = \lim_{F \rightarrow 0+} \frac{(L - L_0)/L_0}{F/A},$$

and then the cantilevered rod's deflection would be given by $\frac{64E'FL^3}{3\pi D^4}$. Then to get perfect rigidity, all we would need is for E' to be zero. There is no reason to think that E rather than E' is the natural magnitude, and to think that both are natural magnitudes offends against parsimony. Moreover, if both are magnitudes, then by the same token we should think that both density and the reciprocal of density are a natural magnitude, which would absurdly forbid a region of space to be empty, as the reciprocal of the density of empty space would be infinity.⁵

3.4.3.3 Immaterial minds In connection with Thomson's Lamp, I already mentioned the possibility of an immaterial mind toggling a physical switch. But one can also exemplify a number of our paradoxes by means of telepathic immaterial minds alone. For instance, while our initial embodiment of the Grim Reaper paradox involved a lamp being turned on, we could also replace the lamp by a mind and the Grim Reapers by other minds each of which is capable of giving the first mind an idea that the first mind could never get on its own. Many of the present paradoxes can be run like that. Huemer, on the other hand, either has to rule out the very possibility of interacting immaterial minds or argue that when these minds are arranged as they would need to

⁵ I am particularly grateful to Ian Slorach for comments on multiple versions of my discussion of Smullyan's rod.

be for one of the paradoxes there would still be infinite intensive magnitudes. To argue for a thesis as controversial as the impossibility of immaterial minds on the grounds of paradoxes of infinity seems to be an overreach, especially given the alternative of causal finitism which has no problems with immaterial minds but rules out the paradoxical infinite causal arrangements of such minds. Perhaps one could indeed argue that immaterial minds arranged paradoxically would have infinite intensive magnitudes. But that is a difficult task.

3.4.4 EVALUATION

Huemer's clever resolution is sometimes unstable under minor modifications to the paradoxes. Furthermore, it is not clear that the magnitudes he needs to be intensive are sufficiently natural. And there appear to be magnitudes that are intensive but that can plausibly be infinite. Some of these magnitudes can perhaps be argued not to be natural, but it is difficult to simultaneously argue them to be non-natural and argue that the ones Huemer needs to be natural are natural. Causal finitism is a superior solution.

3.5 *No room*

Another approach to the paradoxes is to deny that there could be *room* in spacetime for the entities posited by the causal paradoxes of infinity.

A number of the paradoxes make use of an infinite past. So, first, we need:

- (2) It is impossible for time to go infinitely far back.

This by itself is insufficient, since all the paradoxes that make use of an infinite past can also be run in a finite past using a supertask. To that end, we need a discreteness thesis:

- (3) It is impossible for there to be infinitely many instants between two instants of time.

We also need to rule out either temporally backwards causation or an infinite future. For if there can be an infinite future and temporally backwards causation is possible, then we can simply run our paradoxes temporally backwards. For instance, the sequential die-guessing paradoxes in Chapter 5 can be run on the assumption that you perceive the future rolls but not the past ones and must guess the immediate past one. Thus, we need:

- (4) It is impossible for causation to run from the future to the past.

But the multipersonal infinite die-guessing paradox in Chapter 5, Section 2.5 does not require an infinity of causes strung out in time—an infinity of causes all at one time will do. Moreover, it is plausible that simultaneous causation is possible.

Furthermore, if simultaneous causation is possible, then it should be possible for at least some of the paradoxes that involved supertasks or infinite pasts to use a spatial

arrangement of causes. For instance, take the Grim Reaper. Suppose that there is a lamp and a sequence of Grim Reapers each capable of instantaneously and remotely lighting the lamp. The Reapers are arranged left-to-right in a line, at coordinates $\dots, 1/4, 1/2, 1$, or maybe $\dots, -3, -2, -1$. Each Reaper activates at t_0 and checks if any of the Reapers to its left is lighting the lamp at t_0 . If none are, it lights it. These paradoxes may be a little less impressive than their diachronic versions, but nonetheless they have force.

Finally, it is intuitively plausible that if time must be discrete, so must space. Intuitively, if space is continuous, one should be able to move continuously through space, but that would require time to be continuous rather than discrete.

All these thoughts make it very likely that the no-room resolution of the paradoxes will also be committed to:

- (5) Necessarily, space is discrete and finite in extent.

But supposing space to be discrete and finite in extent only helps to avoid the paradoxes involving simultaneous infinities if we suppose:

- (6) Necessarily, all causes are in space.

Otherwise, one could run the paradoxes using non-spatial causes, for instance disembodied minds able to make indeterministic mental events happen and thus able to run die-rolling paradoxes. Furthermore, we need to suppose something like:

- (7) Two distinct fundamental causally efficacious entities cannot be present at the same spatiotemporal location.⁶

For if we have no in principle reason to deny causal infinities, and we allow two distinct causally efficacious fundamental entities to be colocated, it would be *ad hoc* to disallow an infinite number. And if there could be an infinite number, then we could run causal paradoxes of infinity using colocated causes (say, ghosts all present in the same place at the same time).

Thus, for the no-room answer to do the same work that causal finitism does requires quite a large number of theses: (2)–(7). Causal finitism is a much simpler posit. Moreover, some reflection shows that we had to posit all these theses precisely in order to rule out infinite causal histories, i.e., to establish causal finitism. It is preferable simply to accept causal finitism by itself, unless we have independent arguments for the above theses.

Granted, there are some independent arguments for some of the theses. It may be that Zeno's paradoxes support discreteness theses about space and time, for instance. On the other hand, however, there is some *empirical* reason to deny (7). For photons and other bosons can occupy the same quantum state and hence be colocated. And

⁶ The restriction to fundamental entities is to avoid counterexamples such as the lump and the statue made from the lump both of which might count as causally efficacious—say, gravitationally.

even if the ultimate physics turns out to posit fields rather than particles,⁷ it seems very likely that there can be more than one field in a given location: it seems possible for a universe to be suffused both by an electromagnetic and a gravitational field.

We have, thus, not found a satisfactory alternative to causal finitism.

4. Why is Causal Finitism True?

4.1 *The question*

If I am right, then causal finitism is true. But *why* is it true? What “metaphysical force” prevents infinitely many causes from congregating in the causal history of some event? After all, surely, for any finite number n , it is possible to have n causes working together. Why only for finite n ? The explanatory question differs from the justificatory one. If the arguments so far succeed, then we are justified in thinking only finite numbers of items are possible in causal histories, but we do not have an explanation of why there is such a restriction.

The explanatory question is deeply philosophically interesting. But we need not answer it here. Imagine a metaphysician has collated a large number of arguments for the necessary truth of Platonism. It would be really nice if she could give us an explanation of why Platonism is true that goes beyond the claim that it is necessarily true or that such-and-such just is the nature of predication. The desired explanation would, presumably, explain why it is that Platonic entities like properties and numbers exist. It would be an exciting philosophical development. But the inability to give such an explanation would likely do very little to challenge the philosopher’s arguments for Platonism.

It is likewise not necessary to give an explanation of why causal finitism is true to uphold the arguments in the book.

4.2 *Some explanatory suggestions*

Every explanatory project stops somewhere. The main explanatory project of this book stops with causal finitism. I do not here defend a theory as to the metaphysical “force” preventing infinite causal histories. But I can briefly discuss two speculative answers.

One option for explaining why causal finitism is true is given by the no-room story encapsulated by theses (2)–(7). The lack of room is then the source of the metaphysical “force” preventing infinite causal cooperation.

But this is not a unified explanation. The simple and elegant thesis of causal finitism is being explained by a large number of theses about three different topics—time, space, and causation. Moreover, at least one of the theses, namely the denial of

⁷ I am grateful to Ian Slorach for raising this worry.

colocation of fundamental causally efficacious entities, is probably false, as we saw. And we would have a new set of explanatory questions: Why are (2)–(7) true?

A second option would search for an explanation of causal finitism in the nature of the causal relation. From the result of the Appendix in Chapter 2, we can see that causal finitism is the conjunction of two theses: infinite causal cooperation is impossible and causal regresses are impossible. But now suppose that the relevant causal relation—say, partial causation, causal contribution, or causal influence—turns out to be transitive. Then every item in a causal regress leading up to an effect e will itself stand in that causal relation to e .

We can then say that just as it is in the nature of, say, partial causation that it be transitive, it is also in the nature of *being partially caused* that a thing can only have that relation to finitely many others. Of course that raises another explanatory question: Why is the nature of being partially caused like that? But one must stop somewhere in the explanatory project if one is to write a finite book. And thus there is room for future exploration.

5. Further Extensions

5.1 Causal loops

As we saw in Chapters 2 and 3, there are suggestive parallels between infinite causal arrangements and causal loops. It would be very nice if one could unify the rejection of causal loops with the rejection of infinite causal arrangements, and indeed it would provide additional evidence for both theses. We can effect such a unification.⁸

To be concrete, suppose a time-traveling chicken laid an egg in the past which developed into that very chicken. Then (causally) before the chicken, there was an egg, before which there was a chicken, which was preceded by an egg, and then by a chicken, and so on *ad infinitum*. That sounds very much like the kind of causal regress that causal finitism denies. But it's not exactly the same, because these infinitely many causal claims are all about only one chicken and one egg, and that involves two things rather than infinitely many.

What we can do at this point, however, is formulate a graph-theoretic generalization of both causal finitism and the denial of causal loops. Think of a causal nexus as a plurality of nodes—the relata of causal relations—together with directed lines between the nodes, where the nexus contains the directed line $x \rightarrow y$ provided that x is at least a partial cause of y . (I am neither assuming nor precluding that partial causation is transitive.)

If x_1, \dots, x_n is a finite sequence of nodes such that $x_i \rightarrow x_{i+1}$ is a directed line in a causal nexus, then let's say that x_1, \dots, x_n is a *monotonic sequence culminating with x_n* .

⁸ *Compare also how the Axiom of Regularity (see Chapter 2, Section 6) not only rules out backwards set-membership regresses but also membership loops.

Sometimes for clarity I will write the sequence $x_1 \longrightarrow \cdots \longrightarrow x_n$. Note that if there are causal loops, there is no guarantee that the x_i are distinct. We can now formulate a unified thesis that subsumes both causal finitism and the denial of causal loops:

- (8) No possible causal nexus contains a node y which is the culmination of infinitely many monotonic sequences.

Here are four paradigm examples of violations of (8) (Fig. 7.2).

Self-causation: If c is self-caused, then c culminates the repeated monotonic sequences $c \longrightarrow c$, $c \longrightarrow c \longrightarrow c$, $c \longrightarrow c \longrightarrow c \longrightarrow c$, and so on.

Alternation: In the time-traveling chicken (c) and egg (e) causal loop, the chicken culminates the following repeated monotonic sequences: $e \longrightarrow c$, $c \longrightarrow e \longrightarrow c$, $e \longrightarrow c \longrightarrow e \longrightarrow c$, and so on.

Regress: There is a node y that culminates an infinite number of monotonic sequences of the following form: $y_{-1} \longrightarrow y$, $y_{-2} \longrightarrow y_{-1} \longrightarrow y$, $y_{-3} \longrightarrow y_{-2} \longrightarrow y_{-1} \longrightarrow y$, and so on.

Immediate cooperation: A single node y has infinitely many distinct directed lines pointing to it: $x_1 \longrightarrow y$, $x_2 \longrightarrow y$, $x_3 \longrightarrow y$, and so on. The interpersonal die-guessing paradox of Chapter 5, Section 2.5 has this structure.

We can take the arrows here to be whatever relations we decide are the ones to define causal finitism in terms of.

The primary topic of this book is not time travel or causal loops, but infinity. Hence, the extension of causal finitism to ruling out causal loops is something I leave as optional. Nonetheless, the tightness of the analogies between ruling out infinite causal histories and causal loops is very suggestive, and ruling out each strengthens the case for ruling out the other.

We might also prefer to talk about causal priority instead of partial causation. I leave it to the reader to make the requisite changes.

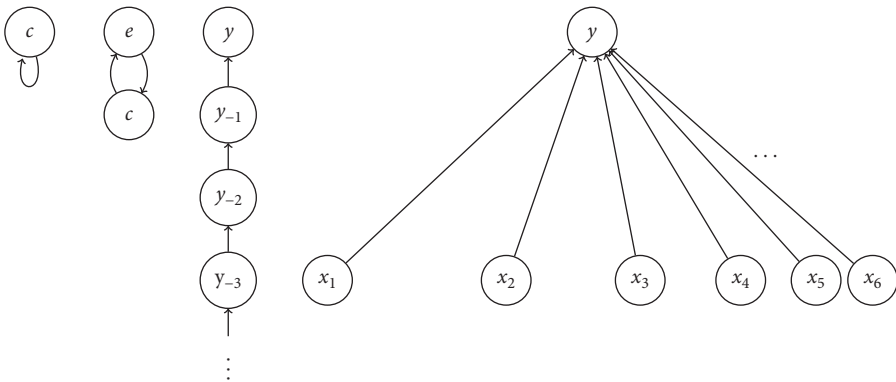


Fig. 7.2 Four paradigmatic violations of (8).

5.2 Explanatory relations

Aristotle (1984, p. *Physics* II.3) famously distinguished four kinds of causes: efficient, material, formal, and final. Each of these corresponded to a different kind of explanation. The causal finitism of this book concerns efficient causation, though in Section 2.2 we briefly considered an extension to material causation in the guise of a principle that no object can be a fusion of infinitely many parts. More generally, material causal finitism could be a mereological finitism, a denial of the possibility of something's having infinitely many parts;⁹ formal causal finitism would deny the possibility of infinite proper definitions; and final causal finitism would deny the possibility of an action or event having infinitely many goals.

And we could generalize even further and simply say that no thing or proposition can have infinitely many partial explanations, and thus get a causal finitism not only for Aristotle's four causes, but also for any further explanatory relations he may have neglected.

There are a number of difficulties with such a complete generalization. Here are four somewhat representative ones.

First, consider some universal arithmetical statement, namely that all natural numbers have some property *P*, say, the property of being even or odd. It seems that this universal truth is grounded in an infinite number of particular truths: 0 has *P*, 1 has *P*, 2 has *P*, and so on.

Second, assuming finitism is false (as I have argued in Chapter 1), this scenario is also possible: there are infinitely many horses and no horse is purple. Then the proposition that no horse is purple has infinitely many propositions of the form $\langle h_i \text{ is not purple} \rangle$ partially explaining it.

Third, someone who believes that she will live forever can be rationally motivated by promoting her well-being tomorrow, the day after tomorrow, and so on, thereby having infinitely many final causes for her action.

Fourth, suppose four-dimensionalism is true and imagine a tree that grows a new branch every day for eternity. The tree, considered as a four-dimensional object, has infinitely many branches as parts, then.

These difficulties are not insurmountable. In the case of natural numbers all having *P*, we can deny that a genuine explanation has been given. Perhaps a genuine explanation of why every natural number is less than some prime number or other would involve giving a finite proof of this fact, and if *P* is a property that for Gödelian reasons cannot be proved to be had by all natural numbers, then there just is no explanation of why every number has *P*.¹⁰

The case of the non-purple horses is more difficult. We might argue, however, that a partial explanation is a part of a full explanation. But the only way to make the

⁹ Cf. Aristotle (1934, p. Book VI).

¹⁰ This suggestion would violate versions of the Principle of Sufficient Reason (PSR) that claim that all truths have explanations. But some philosophical defenses of the Principle of Sufficient Reason apply only to contingent propositions (e.g., Pruss 2006).

propositions (h_i is not purple) be parts of a full explanation is by adding the infinite proposition that every horse is one of h_1, h_2, \dots , and perhaps infinite propositions are impossible, being composed of infinitely many constituents. Thus, no individual one of the explanations is even a partial explanation.

The case of future well-being is perhaps easiest to handle: we do not in fact engage in infinitely many future-directed motivating acts of thought, and so the agent is likely motivated by a single act of thought that universally quantifies over future days. We can describe this in terms of the possession of infinitely many final causes, but that description does not in fact yield a correct explanation of the action, the correct explanation being much more unified. (In Chapter 9, Section 3.2 we will also discuss the more difficult case of God's reasons.)

The infinite tree is difficult. One option is to say that in the case of organic wholes, the whole is entirely explanatorily prior to the parts and the parts are not explanatorily prior to the whole. Thus, perhaps, while explanatory finitism should rule out infinite fusions, it should not rule out infinite organic wholes. Another option is to go for some variety of three-dimensionalism.

Nonetheless, it is far from clear that all variants of such objections to explanatory finitism can be answered. For instance, in Chapter 9, Section 3.3, we will see that theists have reason to reject explanatory finitism for certain kinds of constitutive explanations. Moreover, it is not clear that explanation in general is a sufficiently unified genus to justify a generalization from causal finitism to explanatory finitism.

6. Overall Evaluation

Causal finitism holds, roughly, that causal histories are finite. Some of our choice points relate to the causal relation that generates causal histories, with many options available, such as: full causation, actual partial or contributory causation, actual or contributory ensuring, causally-grounded counterfactual dependence, potential partial or contributory causation, impingement, and unanalyzed causal priority. Of these, actual or contributory ensuring and causally-grounded counterfactual dependence appear to be the best ones overall, though there is something to be said for the potential variety as well. If we go for actual partial or contributory causation, on the other hand, then our overall argument is weakened by an apparent inability to account for the Benardete's Boards paradox, and some versions of the Grim Reaper paradox (ones where there are no fingerprints, no gravitational waves, etc.); still, there are many other paradoxes that the theory can handle.

All in all, here are choice points that we have identified for causal finitism:

- (9) Do we allow for fine-grained events as causes and, if so, do we modify causal finitism to focus on fundamental causes or rule out infinities of fine-grained causes in another way?

- (10) Do we restrict the causal relations that causal finitism speaks of to fundamental ones?
- (11) Does causal finitism apply only to full causation or also to partial causation or causal contribution?
- (12) Do we need to weaken the causal relation further so as to make causal finitism kill the Grim Reaper and Benardete's Boards paradoxes, and if so, how?
- (13) Do we rule out infinite fusions in addition to infinite causal chains or do we carefully analyze the relevant causal relation in order to rule out counterexamples coming from such fusions?
- (14) Does causal finitism apply to causation by absences (variant: privations)?
- (15) Does causal finitism apply to causation of an absence (variant: a privation)?

There is much room for future research here. Causal finitism is a family of theories. It is very plausible that some theory in this family accounts for the causal paradoxes of infinity better than the competitors. The general arguments for causal finitism then give us reason to think that some theory in this family is true, but identifying that theory precisely is a problem I encourage sympathetic readers to work on.

Furthermore, extensions of causal finitism are a promising avenue for exploration. An extension that rules out causal loops is very natural. Extensions to other forms of explanation are more difficult.

Moreover, even with all the room for future research, causal finitism has interesting and non-trivial implications that we shall explore in Chapters 8 and 9. In these chapters I will revert to the rougher formulations of causal finitism that occurred earlier in the book, which will illustrate the fact that most of the applications are independent of many of the details of the questions of how the thesis is to be refined.

8

Discrete Time and Space

1. Introduction

Plausibly, your existence at earlier times causes your existence at later times. Thus, for every instant of time over the past year, your existence at that instant caused your present existence. If causal finitism is true, it follows that there were only finitely many instants of time over the past year. The argument generalizes and so time is discrete—between any two points of time, there are only finitely many points of time.

In this chapter, we will examine this kind of argument from causal finitism to the discreteness of time, its consequences, and its spatial analogue, as well as the nature of the discretism involved. Of particular interest will be whether the argument gives us a *reductio ad absurdum* against causal finitism, on the grounds that time or space could be continuous. However, we will see that the argument from causal finitism to discreteness is not quite as solid as at first it seems. We will show that it is possible to distinguish causal discretism from spatial discretism, and to reconcile this distinction with physics. Causal finitism does provide some evidence for a discrete time and space, but no more than that.

2. Causal Finitism and Discreteness

2.1 *The basic argument*

If an object is falling, it is plausible that for each past time t in its fall, the object's being where it was at t , with whatever velocity it had then, causes it to be where it is now. Examples can be multiplied. Indeed, whether in the context of classical mechanics, Relativity Theory, or Quantum Mechanics, it is very plausible that the present state of the universe is caused by its past state at each past time. Even if the past is finite, as long as time is continuous in the weak sense that between any two times there is another time, it follows that causal finitism is violated. Hence, causal finitism plus a physics like ours seems to require us to take time to be discrete. And what's worse for causal finitism is that standard formulations of the three physical theories all involve continuous time.

2.2 *From discrete time to discrete space?*

Furthermore, if time is discrete, it is very plausible that space is as well. For if time is discrete while space is not, then objects could not move continuously through space.

They would have to make leaps, being in one place at one time, and then at the next discrete time being at a place not adjacent to the original place. But it is very plausible that objects sometimes move through space without hopping like that.

3. Two Kinds of Discreteness

3.1 *Subdivisibility and fixity*

Aristotle held that time is not infinitely subdivided but that it is arbitrarily subdivisible.¹ For any two moments t_0 and t_1 of time (for Aristotle, these would be termini of movement or change), even if there actually isn't a time strictly between t_0 and t_1 , there *could have been* a time there. This is the first kind of discreteness view: time is discrete but arbitrarily subdivisible.

Why might one think this? First, the view does justice to the ordinary intuition that any interval of times can be subdivided. Second, Aristotle viewed time as the measure of change. On a plausible interpretation, instants are the termini of changes. Even if in fact there is no terminus of change between t_0 and t_1 , something *could* have begun or completed changing after t_0 and before t_1 . If that were to happen, a terminus of its change would have constituted a time between t_0 and t_1 .

A second view of discrete time would say that time is a fixed sequence, and the smallest spacing in that sequence cannot be made smaller. There are, in effect, atoms of duration.

The Aristotelian view naturally goes along with a view on which the discrete time is messy, with uneven spacing between instants dependent on the vagaries of where the termini of change fall. The fixed sequence view naturally goes along with a picture of times as evenly spaced.

It is plausible to speculate that whichever kind of discreteness time has is had by space as well. But we would not expect space to have evenly spaced points as on a grid. A regular grid in space would induce a preferred direction, and it is plausible that space does not have such. Think for instance of the three "special" axes at right angles to each other picked out by a cubical grid. On the other hand, a "messy" discrete space, one where points are scattered unevenly, could lack a preferred direction on a large scale.

The "regular" picture of discreteness of space makes rotation problematic. If space were always a regular grid, an object's geometry would necessarily change if it rotated by a small amount. For instance, suppose space is a cubical grid, and consider an object that occupies the eight corners of a cubical cell in that grid. We could rotate that object by a multiple of 90 degrees about an axis passing through the centers of

¹ Aristotle argues that a finite whole cannot be composed of an infinite number of finite parts (Aristotle 1934, Book VI), and applies this to time. He also holds that every motion can be subdivided and that time is the measure of motion (Aristotle 1957, Book IV), so he has to hold that time is always subdivisible.

two opposed faces in the cube, or by a multiple of 120 degrees about an axis passing through two opposed vertices. But we could not rotate the object by, say, 45 degrees about either axis without severely distorting its size and shape. And it seems that it should be possible for objects to rotate in space about arbitrary axes without changing internal geometry.

This argument against the grid picture of space does, however, run into a problem. Given General Relativity, our spacetime is curved. In a curved spacetime, it may well be impossible for objects to rotate or move without slight changes to distances between their components. The changes are tiny in the case of objects in moderate gravitational fields, and that may be enough to pacify our intuitions about rotation. But if space is a grid, then an object spanning a sufficiently large region—namely, a macroscopic one—can rotate with only minor changes to its overall geometry. So if General Relativity can be reconciled with the rotation intuition by invoking approximation, so can a grid picture.

On the other hand, in an Aristotelian picture, there is no need to do any reconciling: the messily arranged points of space could rotate along with the object, maintaining internal geometry. This gives us some reason to prefer the Aristotelian picture in the case of space, and indirectly in the case of time as well.

It is also important to make the historical note that although Aristotle thinks the points of space and moments of time or “nows” are discrete, he does not think they constitute the whole of space or time. Rather, he insists that “two points have always a line between them, and two nows a space of time [*kai tôn nun ton chronon*]” (Aristotle 1934, p. Book VI, translator’s intrusions omitted). This could be used to give a second way to distinguish between views of discreteness, independent of the one just discussed: Is the discrete space or time *constituted* by the points or is there something *bounded* by the points as Aristotle thinks? But while the question of whether there are periods of time between nows is important, it is independent of adjudicating causal finitism, since whether or not there are such intervening periods, there is only one of them per pair of successive moments, and hence the number of periods of time is just as finite as the number of moments.

3.2 Refining the Aristotelian picture

3.2.1 AN OBJECTION TO ARISTOTELIAN DISCRETENESS

There is, however, an argument against the Aristotelian version of discreteness. Even given causal finitism, it should be possible for infinitely many things to come into existence during the last year, as long as the things are set up so that they cannot causally cooperate. But on an Aristotelian picture where there is no fixed temporal sequence, it seems quite possible that the infinitely many things could have come into existence at infinitely many different times. But whenever a thing comes into existence, there is a time on the Aristotelian picture. Or imagine an infinite number of mutually non-interacting radioactive atoms, whose decay times have an exponential decay distribution. Then over any non-empty (and non-infinitesimal) period of time

after the initial setup, with probability one infinitely many atoms would decay.² Thus it seems possible for there to *be* infinitely many different times during a finite period, contrary to the discreteness of time, assuming that there is no predefined sequence of time, but times coincide with the termini of events as the Aristotelian thinks.

Of course, an Aristotelian could also use these thought experiments, combined with causal finitism, as a *reductio ad absurdum* of the possibility of an infinity of past objects in any spacetime, even if these past objects are not causally efficacious. This would still be compatible with the possibility of infinitely many objects spread over an infinity of different spacetimes, or a future infinity of objects, or an infinity of atemporal objects and that would be enough to escape the mathematics-based arguments against finitism in Chapter 1, Section 4.3. But it is not clear that this is a very plausible *reductio*.

3.2.2 INTERNAL AND EXTERNAL DISCRETENESS

There is another Aristotelian option, and that is to make a distinction between internal and external discreteness of time. On an Aristotelian metaphysics, events involve substances in one of four ways: a substance comes into existence, a substance ceases to exist, a substance is accidentally changed, or a substance exercises causality (producing any of the three preceding events). On an Aristotelian picture, thus, it is very natural to take the instants of time to be fundamentally attached to particular substances and to correspond to the substances' entry into and exit from existence, the substances' accidental changes, and the substances' exercises of causality. They are thus instants of *internal* time.

Hence, on Aristotelian metaphysics, time can be seen as fundamentally a feature of particular substances rather than of the world as a whole. One can then try to introduce a shared time by correlating the internal times of the substance. For instance, there is a traditional Aristotelian idea that intersubstantial causation is simultaneous. If a substance *s* at time *t* causes a substance *s'* at its internal time *t'* to enter or exit existence or to be accidentally changed, then we can introduce a single external time corresponding to both *t* and *t'*. In effect, we would be stipulating that the internal times *t* and *t'* of the substances *s* and *s'* are simultaneous. If all goes well—in particular, if there is no time travel or backwards causation—the ordering of the internal times of the substances can then be extended to provide an ordering on the external times of the substances.

There is no *a priori* guarantee that things would go well and generate a shared time that fits well with the ordering on the internal times. There may be metaphysically possible worlds where the internal times do not fit together well enough. It might also be the case that occasional misfits, such as an occasional instance of backwards

² *For any non-empty non-infinitesimal interval *I* of times after the initial setup and any atom, there is a non-zero and non-infinitesimal probability that the atom will decay during *I*. Absent interaction, these probabilities are independent. Thus, by the Law of Large Numbers, with unit probability, infinitely many atoms will decay during *I*.

causation, could be accommodated in some way by positing a shared time that harmonizes *most* of the internal times.

But when things do work out, then we will have a sequence of shared, external times. Every (or every typical, if there are occasional misfits) internal time of a substance will correspond to an external time, since internal times correspond to causal interactions which led to the introduction of an external time. Moreover, every external time corresponds to the internal time of some substance or other.

But typically external times do not correspond to the internal times of *all* substances. After all, at any given external time, some substances will have already ceased to exist and others will have not yet come to be. There isn't even a guarantee that an external time will correspond to an internal time in each substance that exists then. If Bob is a substance, there may be an external time T that corresponds to some internal time t of substance Sally but that does not correspond to any internal time of Bob's, even if both Bob and Sally exist omnitemporally. Rather, T may fall strictly between two of Bob's internal times, say t_1 and t_2 . In such a case, it may be that the correct thing to say about what Bob is doing at T is to say that he is changing between his state at t_1 and his state at t_2 . Thus, those respects in which he remains unchanged between t_1 and t_2 can be attributed to him at T . But if, say, he is turning from red to green between t_1 and t_2 , then perhaps the thing to say is that at T he is neither red nor green but *changing from red to green*.³ Alternately, one could perhaps interpolate between Bob's properties at t_1 and t_2 on the basis of answers to counterfactuals like: If Bob had a time corresponding to Sally's time t , what would Bob be like at that time?

There is much detail to be worked out on such a picture of time, but it is plausible that a genuinely Aristotelian story of time as grounded in change is going to have this kind of a shape. Thus, the Aristotelian discretist has an answer to our argument from the apparent possibility of an infinity of objects with termini of changes that do not fit into a finite set of instants. The answer is that it is only *internal* time that is guaranteed to be discrete. External time need not be discrete. But it is internal time that is metaphysically more fundamental: external time is merely a mathematical construction out of internal times, and attributions of properties to a substance at external times that do not correspond to any of the substance's internal times are derivative from attributions of properties at internal times.

In particular, then, on this account a substance enters into causal relations neither as agent nor as patient at external times that do not correspond to internal times. If

³ A difficulty: What if Bob is changing not from one determinate to another, but simply between having a property and not having it? For instance, perhaps Bob is waking up and becoming conscious. In that case, the above picture suggests that between the two instants Bob is neither non-conscious nor conscious, but changing between the two states. But if we say that he is neither non-conscious nor conscious, then we seem to violate the Law of Excluded Middle. This violation can, however, be escaped. For we could say that Bob is neither (non-conscious)-at- T nor conscious-at- T , but instead (in-the-process-of-becoming-conscious)-at- T . The Law of Excluded Middle guarantees that for any C , Bob is either C -at- T or not C -at- T . But being not C -at- T is not the same as being (not C)-at- T .

Bob doesn't have an internal time corresponding to an external time T , then at T Bob doesn't cause anything—not even his future existence. And hence a non-discrete sequence of external times need not violate causal finitism.

4. Physics

4.1 *An objection to causal finitism*

Standard formulations of major physics theories from Newton onward either model time with the real numbers or model spacetime as a continuous manifold with local coordinates corresponding to quadruples of real numbers.

Moreover, this is not merely an accidental feature of the theories. The continuity involved is essential to the differential equations in which the laws of physics are couched. One could, of course, produce a theory based on a discrete space and time that produces empirical consequences so similar to Newton's theory that we could do no experiment to tell them apart. But the discrete formulation would be apt to be much more complicated. And simplicity is always a part of the attractiveness of the major theories.

This supports the second premise of the following major objection to causal finitism:

- (1) If causal finitism is true, time is discrete.
- (2) Time is not discrete.
- (3) So, causal finitism is not true.

Thinking through this argument will lead the causal finitist to two options. The first will be to embrace a speculative physics that is not committed to the continuity of time or spacetime, thereby questioning the grounds for (2). While much of physics from Newton to recent times has supposed time and space to be continuous, there are now physically live discretist options. For an excellent survey of options, see Hagar (2014). I have nothing to add to this.

I will, however, explore a second approach, which is to interpret the causal import of quantum physics in a way that is compatible with causal finitism but does not substantially change the physics of the theory. The result accepts (2) but denies (1).

4.2 *Causation and physics*

Formal formulations of typical physics theories do not use the word "cause". This might lead one to suppose that the concept of causation is no longer needed. But that would be a mistake. It is not possible, after all, to describe what the experimental physicist does in the lab without causal vocabulary. Buttons are *pushed*, outcomes are *observed*, etc. (cf. Anscombe 1971). But the lack of causal vocabulary in the theories does open options.

When a physical theory describes the evolution over time of a system, it is natural for friends of causation to read causation into this evolution by supposing that the earlier states of the system cause the later ones. This reading of physics is central to the argument from causal finitism to the discreteness of time. But we need not read the physics like that.

It is trivially easy to reconcile continuous time, causal finitism, and a physics that does not use the word “cause”. Begin with a crazy thought: causation happens only once per year. For convenience, I will take a year to be a half-open interval $(n, n + 1]$ of times, which includes a last moment but not a first one, and where I take years to be units of time, and suppose that integers like n or $n + 1$ correspond to the last moments of each year.

For each year $(n, n + 1]$, there is a non-instantaneous state s_n of the universe over that year. We could then suppose that s_n causes s_{n+1} which causes s_{n+2} and so on. And there is no causation other than between these year-long states.

Alternately, we could suppose that at each time t there is an instantaneous state of the universe, u_t . Moreover, most instantaneous states u_t of the universe are causally inert, with the exception where t is an integer n . Thus, the only causally efficacious states are the last instantaneous states in each year. We then suppose that either (a) each such instantaneous end-of-year state u_n directly causes all of the “temporally fat” state s_n of the universe over the succeeding year $(n, n + 1]$, including the end-point state u_{n+2} , or (b) each end-of-year state u_n , where n is an integer, directly causes each of the infinitely many instantaneous states u_t for t in $(n, n + 1]$ (with the causation from u_n to u_t crossing a temporal gap of length $t - n$).

Both of the stories involve causation across temporal distance. The state of the universe over the year $(2016, 2017]$ or at the very end of that year causes not only the state of the universe at the beginning of January 1, 2017, but it also causes the state of the universe in March 2017. And it does not cause the state of the universe in March *by* causing states in January and February—it does so directly.

Since the physics theories do not talk of causation, this does not affect the theories or their empirical import. However, positing one instance of causation per year would also undercut the argument that making sense of what the experimental physicist does in the lab requires positing causation. We will do better with respect to common sense if we suppose the causal “ticks” to be not annual but on a finer scale, particularly a scale finer than that of the human sense of time.

Furthermore, there is an obvious common-sense objection to the one causal tick per year theory. We causally explain the state of the universe in July by means of the state of the universe in June. But on the story as given, the states in June and July are simply common effects of the previous year or by its last moment, and there is no causation from June to July. But it gets worse. Not only is there no causation between June and July, but there is no explanation either. One can make *predictions* about July on the basis of June, but the real explanatory work is all done by the preceding year or its last moment.

This objection can be adapted no matter how short the ticks between instances of causation, but it is less intuitively compelling when the ticks are short. It would, indeed, be a significant cost, and far too reminiscent of occasionalism, to deny that the rainy June explained the high water levels in July, and to say that both are caused by the events of the previous year. But as the ticks get shorter and shorter—especially if they should get to the Planck time—the intuitive cost is much smaller. We can say that the rainy June explained the high water levels in July, though we cannot say that the state of the world at a quarter of a Planck time after a tick explains the state of the world half a Planck time after it. The latter is counterintuitive, but we have already from Quantum Mechanics that reality is strange at small spatial scales, and strangeness at small temporal scales should not be a big surprise.

The lesson learned from this trivial escape from the argument from causal finitism to discrete time is that causal finitism requires discreteness in the order of causation, but discreteness in the order of causation is logically compatible with continuity in the order of time. Nonetheless, the causal tick approach is inelegant and *ad hoc* in that the causal ticks are not grounded in the physics—they are an arbitrary add-on.

The physics objection to causal finitism has two versions: the stronger says that the actual world's physics involves continuous time and the weaker says that it would be possible to have a physics with continuous time. The *ad hoc* nature of the above escape is not much of a problem given the weaker version of the objection. It should not surprise us if some false physics (say, Newtonian physics), in order to be made causal, would require *ad hoc* metaphysical posits.

Fortunately, the *ad hoc* periodic causal-tick solution is not always the only escape from the argument from causal finitism to discrete time. I will use Quantum Mechanics to illustrate how a non-Newtonian physics might allow one to reconcile continuous time with discrete causation without being *ad hoc*. Then I will offer an interpretation of Quantum Mechanics on which time is discrete in an irregular way, more like on the Aristotelian discrete time theory than the regular sequence theory. It is not my point to describe how things are, but merely to sketch how physics *might* be reconciled with causal finitism. Deciding between the options is a good task for future research.

4.3 Quantum collapse

4.3.1 SOME BACKGROUND

Quantum Mechanics has two central parts. First, there is the Schrödinger equation. This is an equation governing the evolution of the wavefunction over time and insofar as this evolution is governed by the Schrödinger equation, it is completely deterministic: the future states of the wavefunction are entirely determined by the past states. However, despite the determinism at the level of the wavefunction, we have very strong empirical reason to think deterministic wavefunction physics fails to determine the outcomes of observations. Two electrons can be prepared with the same wavefunction, sent through the same magnetic field, and yet observation can

show them going in different directions, and all we can predict from the wavefunction are probabilities of different observations.

Such predictions are linked to a second part of Quantum Mechanics: the Born rule, which gives a mathematical specification of how to get probabilities of particular observations from the wavefunction.

The problem of the interpretation of Quantum Mechanics is how to make these two components work together. We can divide up the interpretations of Quantum Mechanics into four families:

- (i) Keep the Schrödinger equation exceptionless and explain observations solely in terms of the wavefunction.
- (ii) Keep the Schrödinger equation exceptionless but explain observations in terms of physical features of the world that go beyond the wavefunction.
- (iii) Keep the Schrödinger equation exceptionless but explain observations in terms of non-physical features of the world that go beyond the wavefunction.
- (iv) Make indeterministic exceptions to the Schrödinger equation to account for indeterminacy of observation.

The first family of interpretations appears the most parsimonious, and includes as its main instance the Everett (1957) multiple-worlds interpretation. On the Everett interpretation, the wavefunction evolves deterministically, but what it describes is a branching multiverse. In situations where it seems to us that there is underdetermination of the observations by the wavefunction, what actually happens is that the world is split into a branch where one observation is made and another branch where another is made.

There are, however, serious problems on the Everett interpretation with how to make sense of the probabilities in the Born rule. Suppose a physical situation where there are only two possible observations and the Born rule specifies that one of them has probability $2/3$ and the other $1/3$.⁴ The universe, including the experimenter, is split in two upon measurement, with one experimenter observing one outcome and the other the other. Ontologically this is symmetric: there are two observations made, by two experimenters. So how can one of the two observations count as more likely? There is, of course, a literature on this difficult issue (e.g., Greaves 2006), but notwithstanding defenses of the view, it gives us good reason to pursue alternatives.

The best known representative of the second family is Bohm's (1952) deterministic physics, which in addition to the wavefunction posits particles that travel along trajectories determined by the wavefunction. The outcome of observations, then, is not a function of the wavefunction but of the positions of the particles. The physics here is deterministic, but nonetheless one can statistically recover the probabilistic predictions of the Born rule given a special assumption about the statistical properties

⁴ *For instance, the measurement of the spin of a system in state $\sqrt{2/3}|\text{up}\rangle - \sqrt{1/3}|\text{down}\rangle$.

of the initial distribution of the particles. There are some conceptual difficulties about how to make sense of such probabilities. These difficulties appear to be no greater, but also no lesser, than those about how to make sense of probabilities in classical thermodynamics with deterministic Newtonian particle physics.

The main representatives of the third family are dualistic supplementations of the Everett interpretation. There is, as in Everett, a deterministic multiverse. But the problem of the recovery of probabilities from the deterministic multiverse is solved by positing non-physical minds that travel indeterministically through the deterministic multiverse. Thus, when one sets up an experiment that has a $2/3$ probability of leading to one observation and a $1/3$ probability of another, what happens is that there is a $2/3$ chance of a mind traveling up one branch and a $1/3$ chance of a mind traveling up the other branch. These interpretations are further subdivided depending on the number and dynamics of minds. On the many-minds view, there are infinitely many minds, and whenever a branching happens, infinitely many take each branch (Albert and Loewer 1988). So there are infinitely many minds corresponding to the experimenter's body (we might take bodies to be delineated by the wavefunction). On single-mind views, there is at most one mind per body, and the minds either travel through the multiverse independently or are nomically constrained so that all the minds are always in the same branch of the multiverse.⁵

The fourth family is the collapse interpretations. A wavefunction that is compatible with a multiplicity of observations—say, observing an electron at one location or observing it at another—*collapses* into a wavefunction that is compatible with only one observation. The collapse is an indeterministic process whose probabilities are such as to yield predictions that fit with the Born rule. This family divides into two subfamilies, depending on the conditions that trigger collapse. On one subfamily—the von Neumann interpretation—collapse is triggered by observation itself. On the other, it is triggered stochastically or deterministically by some other physical conditions, typically ones correlated with events at levels of energy or size scale beyond those found in typical quantum interactions. The best known member of this subfamily is the Ghirardi–Rimini–Weber (GRW) collapse theory (Ghirardi, Rimini, and Weber 1986).

4.3.2 CAUSATION

The Everett and Bohm interpretations which are the main representatives of families (i) and (ii) do not lend themselves neatly to a reconciliation of causal finitism with

⁵ The standard objection to the independent travel version is that on this view it is likely that the minds have spread out through the multiverse so much that we are unlikely ever to meet up with a mind, and hence the bodies we are meeting up with are predominantly mindless zombies (Albert 1992, p. 130). The constrained version is found in Barrett (1995) and avoids this problem. There is also an Aristotelian variant on which the multiverse is traversed not just by minds but by all forms, including those of non-minded substances, even inanimate ones (Pruss 2018).

continuous time that isn't *ad hoc* in the way my annual causal tick theory was. There just aren't natural transition points for introducing causation.

However, the dualist branching multiverse and the collapse interpretations of (iii) and (iv) have a potential of fitting more nicely with discrete causation. We can suppose that the wavefunction (or, more precisely, the physical reality that the wavefunction describes) and prior state of the minds indeterministically cause the minds to take one or another branch in the dualist multiverse interpretations. And we can suppose that each case of collapse is an instance of causation.

There is then reason to hope that the occurrences of causation will be temporally discrete, at least in a finite universe (or perhaps in a finite subsystem of an infinite universe). Note that in the branching interpretations, all we need is that there be discreteness *within* each branch from the beginning onward.

On the branching multiverse version, we can then suppose that there is no causation within the temporal evolution of the wavefunction. Thus either the wavefunction, considered as a temporally extended entity, is itself an uncaused cause of the travels of the minds or the collapses, or else there is some other cause that causes the whole of the temporally extended wavefunction. It is counterintuitive that the evolution of the wavefunction is acausal, that the wavefunction is caused *en bloc*. Nonetheless, ordinary causal language can be taken to refer to the causation involved in minds being caused to take one branch or another or the wavefunction being caused to collapse.

On the collapse versions, on the other hand, we can take each instance of collapse to cause all of the temporally extended state of the wavefunction until the next collapse. And between collapses, the evolution of the wavefunction will still be acausal. As long as collapses happen often enough, this will preserve intuitions about the causal efficacy of lab work and everyday life.

All of these discrete causality readings of interpretations of Quantum Mechanics have the same empirical consequences as the underlying interpretations. The only additional thing that is done by this reading is the introduction of an account of where causation occurs in the physics.

The above solutions may have difficulties in infinite quantum universes or multiverses. But given that Quantum Mechanics appears to involve instantaneous action at a distance in cases of entanglement, causal finitism may also force quantum systems to be finite in order to avoid causal infinities.

4.3.3 BACK TO DISCRETE TIME

The dualist branching multiverse and collapse interpretations of Quantum Mechanics thus allow for discrete instances of causation interspersed with acausal deterministic evolution of the wavefunction according to Schrödinger's equation. On these interpretations, causation is discretely arranged but time is continuous.

But now it turns out that there is also a way of regaining discrete time if the multiverse or universe is finite. The trick is to take as actual times only those instants of time

at which branchings or collapses occur and to take only the values of the wavefunction at those times to reflect physical reality. On collapse theories we will then have two options: the wavefunction at the collapse time is either the pre-collapsed or post-collapsed wavefunction. Since again my point is merely to sketch how things might be, I will for concreteness take the view that at each collapse time the wavefunction is in the post-collapsed state.

As long as the time of the next branching or collapse and the value of the wavefunction⁶ just before it is a deterministic or stochastic function of the value of the wavefunction at the preceding branching or collapse time, we do not actually need to suppose that the wavefunction's values between branching or collapse times correspond to anything physically real. These values can be just taken to be a mathematical fiction, with the physically significant values of the wavefunction being the ones at the times of collapse. If we like, we can introduce fictitious mathematical times between the real branching or collapse times for the sake of mathematical convenience. If this is attached to GRW theory, it will be a version of Bell's (1987) flash ontology, with a semantics that allows claims to be made about what is happening at non-flash times, namely the wavefunction having values at them.

The resulting picture gives us instants of time that are irregularly spaced. Their spacing can be explained by means of the wavefunction at the instants and Schrödinger's equation. On dualistic branching views, the spacing of the instants can be read off from the wavefunction. This yields an Aristotelian messiness of spacing, but with the instants being fixed unlike in Aristotle. On the collapse view, however, the instants are not fixed. Had previous collapses gone differently, later collapses would be likely to occur at other times, and so there is nothing fixed about the instants, and the picture is even more like Aristotle's.

This is a non-relativistic picture. In any case, it is not known how to make Quantum Mechanics work with General Relativity, and the pictures I offered are only toy models of how one could have discrete causation along with continuous or discrete time.

5. Fields and Discrete Space

Assuming a flat spacetime, the value of the electromagnetic field at a time t and a location x depends on the values of the electromagnetic field at time $t - \Delta t$ in a ball centered at x with radius $\Delta t/c$, where c is the speed of light. Hence, if space is continuous, the value of the field at any point depends on an infinite number of values of the field at any given past time. If this dependence is causal, as it seems to be, we have a violation of causal finitism.

⁶ **The Schrödinger equation is the partial differential equation $i\hbar\partial\Psi(x,t)/\partial t = \hat{H}\Psi(x,t)$ which is first-order in the time coordinate, and so solving it requires only the value of the wavefunction at a time, and no values for time derivatives. To use a similar method with a second-order equation like that in Newton's laws would require attributing physical reality to both values and their time derivatives.

If space is discrete, the problem will disappear as there will only be finitely many points within a ball of finite size. Thus, we have another argument from causal finitism to discrete space. And if space is discrete, it is plausible that so is time.

But let's get a little clearer on the assumptions in the argument. First, the argument requires a realism about the values of the field in empty space. One could think that there are charged objects, magnets, conductors, and the like, but be a non-realist about the electromagnetic field, holding that it's a useful mathematical fiction for giving an account of the action of objects. On this view, objects directly act on each other at a spatial and temporal distance as if there were an electromagnetic field mediating that action.

Second, the argument requires that facts about values of the electromagnetic field in empty space be causally efficacious. Consider a simplified case. Start with empty space and no electromagnetic field. Charged objects A_1 and A_2 then suddenly appear at rest at locations x_1 and x_2 , respectively, at time t , with distance d from x_1 to x_2 . At time $t + d/c$, each object experiences an electrostatic force from the other object. We could be instrumentalists about the electromagnetic field between the objects. But we could also keep the intuition that the field causally mediates between the objects. However, instead of thinking that the field does this bit-by-bit, we could suppose that the objects x_1 and x_2 at time t work together to *directly* cause the whole of the electromagnetic field over the four-dimensional region between time t and time $t + d/c$ (or over the intersection of that region with the union of the light-cones centered on (x_1, t) and (x_2, t)). This electromagnetic field then causes the movement of the objects at time $t + d/c$. The values of the electromagnetic field at intermediate times—say, at $t + d/2c$ —are causally inefficacious. Further development of this account could perhaps marry it with the discrete quantum causation story in Section 4.3.2.

There is another way one can get out of the field argument from causal finitism to discrete space. Suppose we think of the electromagnetic field as a spatially simple extended entity. The field isn't made up of little bits of field here and there. It is a single field through all of space, without parts. It has values in different places, but these values should not be seen as grounded in properties of localized parts of the field. Instead, they are grounded simply in a single highly determinate global distributional property (cf. Parsons 2000) of the field—its global value, which may be mathematically represented by a function from points of space to values, but is nonetheless a single fundamental determinate.⁷ In that case, the causal history of the fact that the field has such-and-such a value at time t need not include an infinite number of causes at time $t - \Delta t$. It can simply depend on the global state of the field at time $t - \Delta t$.⁸

⁷ Compare this: the location of a point-sized object A in Euclidean space can be mathematically represented as a function f from the set $\{1, 2, 3\}$ to real numbers, where $f(n)$ is the value of the n th coordinate. But we should think of the location of A as more fundamental than the values of f at the three points of $\{1, 2, 3\}$: the latter is a mere representation of a location which subsumes all the coordinates.

⁸ For a critical discussion of distributional properties, see McDaniel (2009).

In other words, the field argument from causal finitism to discrete space is not very strong. We can take it as an argument for a large disjunction. If causal finitism is true, then we have discrete space, or non-realism about fields, or causally inefficacious fields in empty space, or globalism about fields.

6. Evaluation

Causal finitism makes it somewhat plausible that time, and perhaps space as well, is discrete. This discreteness can be fixed and regular, fixed and messy, or flexible and likely messy. However, although causal finitism makes this discreteness plausible, it does not force it. It is possible to interpret physics in a way that makes the causal order be discrete even though space and time are actually continuous.

Causal finitism makes discretist physics more attractive, and hence should encourage physicists to explore that option. But it does not *force* discretism about time or space.

In Chapter 7, Section 3.4, I considered Huemer's account of many of the same paradoxes I discussed. One of the problems with Huemer's account was that it placed significant constraints both on current physics (e.g., the rejection of black holes as described by General Relativity) and on any hypothetical physics.

There are two ways of understanding physics. The classical Aristotelian way is that physics is an account of natural causation. If physics is understood in this way, then causal finitism also places significant constraints, especially by requiring causal chains to be discrete—i.e., to have a finite number of links between any two points—even if time does not have to be discrete. The discreteness of causation requirement is one that Aristotle accepted, but it is nonetheless a constraint.

The other way of understanding physics is to see it as a description of the laws describing the evolution of the universe, with the laws being neutral on causality. On this understanding of physics, which seems to fit particularly well with the practice of modern theoretical physics, causation belongs to the philosophical interpretation of physics, and hence the discreteness requirement is a metaphysical constraint not on physics but on the philosophy of physics. Huemer's constraint, on the other hand, was a philosophical constraint on physics itself, on either understanding of physics. It is more reasonable for metaphysics to constrain the *philosophy* of physics.

At the same time, it is hard to see a plausible causal interpretation of a cosmology with an infinite past history that would not violate causal finitism. Thus, causal finitism does constrain cosmology by forbidding infinite past histories. This constraint, however, does not require any changes to our current best cosmology, since that cosmology assigns a finite age to our universe (starting with a Big Bang under 14 billion years ago), whereas Huemer rejected the current best physical description of black holes.

9

A First Cause

1. Introduction

We begin by showing that causal finitism very quickly implies the existence of a first cause, and then I argue more controversially that there is good reason to take this first cause to be a necessary being. This results in a cosmological argument similar to the Kalām argument.

The most prominent theory of a necessarily existing first cause being theism, I then consider the coherence of the theory of this book with theism. For while many theists will be gratified by causal finitism providing a vindication of a cosmological argument, there is a tension between theism and the arguments of this book. First, we consider an objection from Jonathan Kvanvig that causal finitism contradicts the plausible thesis that God makes decisions on the basis of infinitely many factors. Considerations about the relation between thoughts and their contents as well as classical theism's doctrine of divine simplicity come to the rescue. Second, we consider the causation between contingent events and God's knowledge of them, which threatens to violate causal finitism if there are infinitely many contingent events. Classical theism, however, motivates one to reject the idea that the relationship between the events and God's knowledge is causal. Third, a number of the arguments of the book have relied on the idea that certain pieces of knowledge—say, about how infinitely many die rolls turned out—require a causal connection. But then either God can know these things or he cannot. If he cannot, then omniscience is violated. If he can, then omniscience is saved, but either causal finitism is violated or there seems to be the possibility of God communicating his knowledge in a way that subverts some of the arguments for causal finitism. To solve this problem, it is argued that there is reason to extend causal finitism to what one might call quasi-causal finitism.

2. An Uncaused Cause

2.1 *The quick argument*

There is a quick argument from causal finitism to a first cause:

- (1) Nothing has an infinite causal history.
- (2) There are no causal loops.

- (3) Something has a cause.
- (4) Therefore, there is an uncaused cause.

To see that this is valid, suppose that a_0 is something that has a cause, say a_1 . For a *reductio*, suppose that every cause is caused. Since every cause is caused, a_1 has a cause a_2 , which has a cause a_3 , and so on. Moreover, by (2), the a_i are all distinct, which implies that a_0 has an infinite causal history, contrary to (1). So there must be an uncaused cause.¹

In fact, the argument establishes something a little stronger than that there is at least one uncaused cause. It shows that every caused item a_0 has an uncaused item—a first cause of a_0 —in its causal history. (Though note that if causation turns out not to be transitive, then a “first cause” of a_0 need not actually be a cause of a_0 ; it just needs to be at the beginning of a causal chain leading to a_0 .) Moreover, every backwards causal sequence starting with a_0 will go back to some first cause of a_0 . The plurality of all uncaused items then has the property that every causal chain goes back to some member of that plurality.

2.2 Towards a necessary being

All the premises of our cosmological argument other than causal finitism (i.e., (1)) are intuitively very plausible. Moreover, there is no Causal Principle or Principle of Sufficient Reason among these premises, as in many cosmological arguments (e.g., see Pruss 2012). The reason no such principle is among the premises is that our conclusion is modest: there is an uncaused cause. We have seen that we can add that such a cause is found at the head of every causal chain, but even this conclusion is quite compatible with some or all of these uncaused causes being contingent, whether they be minor items like an uncaused brick or major ones like the Big Bang. This is far from an argument for the existence of God.

But the argument would become more impressive if we could add a Causal Principle like the following:

- (5) Every contingent item has a cause.²

For then any uncaused cause would have to exist necessarily, and we have the more impressive claim that there is a necessary being at the head of every causal chain. Of course, even this claim falls significantly short of theism. For instance, the claim is compatible with the first cause being a necessarily occurring Big Bang, or there being

¹ **The argument as worded here uses the Axiom of Dependent Choice since for each a_i , we need to choose an a_{i+1} . However, we can make use of (1) to avoid any use of Choice. Since the causal history of a_0 is finite, there is a one-to-one map ϕ between the members of that causal history and the integers $\{1, \dots, n\}$ for some n . We can then replace the arbitrary choice of a_{i+1} with a specified choice of the cause c of a_i that has the smallest value $\phi(c)$.

² If items that come into existence in time must be contingent, this is stronger than the Causal Principle in Chapter 3, Section 3.6.

some larger plurality of physical events or entities that are at the head of every causal chain. Nonetheless, the existence of a causally efficacious necessary being would itself be quite an interesting conclusion,³ and further research would be needed to figure out the nature of such a being or such beings.

2.3 Support for the Causal Principle

The Causal Principle (5) is intuitively quite plausible, plausible enough that we should accept it absent defeaters. We already considered the main defeaters to a different version of the Causal Principle in Chapter 3, Section 3.6.2, and the defeaters there are also the main ones available for the present Causal Principle.

It is also interesting to note that the Grim Reaper paradox can also be used to support the present Causal Principle. For suppose that contingent items can come into existence without cause. Suppose that we are in a world where time has the following subdivisibility property: for any distinct times t_1 and t_2 , there is a *possible* time u strictly between them.⁴ Repeating this line of thought, we conclude⁵ that there is an infinite sequence of possible times u_0, u_1, \dots such that u_0 is strictly between t_1 and t_2 and u_n is strictly between t_1 and u_{n-1} for $n \geq 1$.

Much as in Chapter 3, Section 3, suppose a lamp is off at t_0 , can only be activated by the pressing of a switch, and cannot be deactivated. But whereas previously we imagined Grim Reapers that were set for a particular time, we now imagine Instant Grim Reapers (IGRs) that activate as soon as they come into existence. When an IGR comes into existence by the lamp it checks whether the lamp is on. If it's off, it instantaneously and infallibly lights it. Either way, as soon as it's done, it jumps away from the lamp.

If contingent items can come into existence causelessly, so can IGRs. It would be *ad hoc* to allow some contingent items, but not IGRs, to come into existence causelessly. Moreover, the IGRs should be able to pop into existence at any possible time (which would then be actual) after t_1 , and in particular at u_n for any natural number n . Moreover, whether an IGR comes into existence causelessly at one time near the lamp would seem to be completely independent of whether IGRs came into existence near the lamp, as long as the others weren't blocking the space near the lamp—which they are not since they jump away from the lamp when done. Causeless poppings into existence should be independent of one another, barring concerns about room in space.

Given this independence, it should be possible that an IGR pops into existence at u_n for *every* natural number n .⁶ But now we have the classic Grim Reaper paradox once again. The lamp can't be deactivated once activated, and only an IGR would have

³ Pruss and Rasmussen (2018) is a whole book devoted to arguing for this conclusion.

⁴ I won't worry about the details of the ontology of possible times. Perhaps possible times are given by numbers in some temporal measurement system.

⁵ **The Axiom of Dependent Choice is being used here.

⁶ This is similar to the argument for a limited Axiom of Choice in Chapter 6, Section 4.

activated it. But no IGR could have activated it for exactly the same reasons as in the classic paradox: if the IGR that came into existence at u_n turned it on, then the one that came into existence at u_{n+1} (since $u_{n+1} < u_n$) must have failed in its job, whereas IGRs don't fail.

So just as the Grim Reaper paradox supports causal finitism, it also supports the Causal Principle (5).

The most serious difficulty in this argument is the assumption that it is possible for time to be infinitely subdivisible. A second difficulty is with handling, in a non-question-begging way, the possibility that the lamp turns on for no cause at all. However, given the independence of causeless events, it would be *necessary* that if in this scenario an IGR popped into existence at every u_n , then a *further* causeless event would have to have happened, and happened not on account of the IGRs' activities, namely the turning on of the lamp. Furthermore, these uncaused causes would be necessitated to happen by the pattern of IGRs, and that is very hard to swallow.

2.4 The Kalām argument

Our cosmological argument for a necessary first cause has a central feature in common with the Kalām versions of the cosmological argument (Craig 2009), which go back to medieval Islamic philosophy. Like the Kalām argument, our argument denies the existence of backwards-infinite sequences. But the details are quite different. The Kalām argument opposes sequences that are *temporally* backwards-infinite, while the present argument denies all *causally* backwards-infinite sequences.

3. Compatibility with Theism?

3.1 Theism

In Chapter 1, I argued against finitism, and hence for the possibility of an actual infinite. On the other hand, I have now offered a cosmological argument for a necessarily existing first cause. The most prominent theory about a necessary being who is a first cause is that it is God, a necessarily existing perfect being, and so the arguments of this book support the existence of such a being. But are the arguments of this book compatible with the existence of such a being?

3.2 Divine motivation

Consider this argument, suggested to me by Jonathan Kvanvig:

- (6) If God exists, God's creation of the cosmos is made on the basis of infinitely many reasons.
- (7) An action done on the basis of a reason is caused by that reason.
- (8) So, if God exists, causal finitism is false.

There are at least two paths to (6). The first path is to think of the infinitely many possible scenarios as actually competing for actualization.⁷ We should not take the reasons favoring scenarios other than the ones God actualized to be reasons on the basis of which God chose as he did. Those reasons militated against what he did and did not causally contribute to his actual decision.⁸ They counter-contributed, instead, and it may be that causal finitism does not need to rule out infinitely many counter-contributions.

The second path is one suggested to me by Jonathan Kvanvig. There are infinitely many reasons favoring the actualization of our own world. But divine omnirationality implies that God acts on all the reasons that favor the actualization of our world.⁹ And that yields (6).

It is particularly plausible that there are infinitely many reasons favoring the actualization of our world if there is an infinite number of goods g_1, g_2, \dots in the world, say an infinity of good future days. For then the fact that the world includes good g_n is a reason in favor of actualizing this world.

It is difficult, thus, to deny (6). I shall instead deny (7), along a line already sketched in Chapter 5, Section 3.2. We can understand a reason as a mental content or a thinkable favoring an action. A reason is thus something abstract. But in addition to the mental contents or thinkables, there are the token thinkings that realize these contents. It is not the reasons considered as abstract thinkables that are causes of an agent's actions. Rather, it is the token thinkings that realize these thinkables that are the causes of an agent's actions.

The arguments in favor of (6) made it plausible that there are infinitely many thinkables on the basis of which God created as he did. However, (7) is only plausible if reasons are taken to be the token thinkings of the thinkables.

Can we argue for (6) while consistently taking "reasons" to be the token mental acts? It's doubtful. Multiple thinkables can be realized in a single act of thinking. As noted in Chapter 5, Section 3.2, when one believes the moon is round and gray, one thereby also believes that it is round and that it is gray. Likewise, multiple reasons can be realized in a single act of thinking.

Furthermore, the doctrine of divine simplicity implies that *all* of God's thoughts are found in a single act. The solution of relating all the thinkables to one act of thought is not arbitrary, but an integral part of classical theism.

Besides this, it is not entirely clear whether we should think of the relationship between divine actions and divine reasons as a causal one.

⁷ Leibniz says that each possible world has "the right to lay claim to existence to the extent of the perfection it enfolds" (Rescher 2013, Section 54).

⁸ A similar argument was used in Chapter 5, Section 3.2.

⁹ Or at least all the unexcluded ones. See Pruss (2013b) for more discussion. But it is unlikely that all but finitely many are excluded.

3.3 *Divine knowledge*

It seems possible for there to be an infinite number of contingent events not determined by God, even if causal finitism is true. For instance, perhaps some or all people will live forever (e.g., in an afterlife) and will continually make undetermined free choices. As long as each choice depends only on a finite past sequence, there need be no violation of causal finitism. But if an omniscient God exists, he will know the outcomes of these infinitely many choices. Furthermore, it is plausible that no one, not even God, can determine free choices. Thus, God's knowledge must be a kind of reaction to the free choices, and hence an infinite number of free choices causes a single event, say God's knowing the conjunction of the propositions reporting the events.

There are a number of possible responses to this argument. Open theists deny that God knows future free actions. One could combine this with the doctrine that there cannot be any past-or-present infinities (cf. Chapter 7, Section 3.3), and so at any given time God only knows finitely many free actions as only finitely many free actions have occurred at any moment.

On the opposite side of the theological spectrum, we have theological compatibilists such as Calvinists and some Thomists who hold that God causally determines all contingent items, including the free choices of agents. On this view, God could know an infinite number of contingent events by making, and knowing he is making, a single efficacious decision to bring about all the infinitude of contingent events. God's decision might be causally prior to his knowledge, but that need not violate causal finitism.

But these are extreme and problematic views. Open theism rejects the doctrine of God's omniscience as that has been traditionally understood in the Jewish, Christian, and Islamic traditions. Moreover, by ordinary inductive reasoning, we fallibly know about some future free actions of agents. One can know on the basis of a person's past performance that she will freely choose to keep a promise, say. Then, given that open theism leaves no room for God to infallibly know such future truths, we are left with two problematic options (cf. Kvanvig 1996). Either God has fallible knowledge or there are some things that we (fallibly) know which God doesn't know at all, whether fallibly or not. If God has fallible knowledge, then by the same token he has fallible beliefs, and hence it is possible for God to be wrong, which is deeply problematic. But it is also deeply problematic if we know some things that God doesn't.

On the other hand, theological compatibilism runs into serious difficulties with the problem of evil. For if God can determine free choices, then God can make us all *freely* choose the right thing, and it is particularly difficult to justify his permitting evil. Moreover, if our choices are determined by God, then it appears that God determines us to do evil sometimes, and that in itself appears incompatible with moral perfection.

Of course, there are responses by open theists (e.g., Taliaferro 1993) and theological compatibilists (e.g., a number of the authors in Alexander and Johnson 2016) to such

objections, and it goes beyond the scope of this book to debate these. Nonetheless, the arguments are compelling enough that we need to take seriously the middle of the road position that God does not determine free actions and yet has exhaustive knowledge of what will and will not be freely done.

On this middle of the road position, the infinities of future free actions are explanatorily prior to God's knowledge of them. But even so, this is only a violation of causal finitism provided that this explanatory priority is causal in nature. However, it need not be. There is independent reason in classical theism to deny that creaturely actions can cause divine mental states: such creature-to-God causation appears to violate divine aseity.

But if the explanatory relation between creaturely free action (and perhaps creaturely stochastic events) and God isn't causal, what is it? That is a difficult question, and one that takes us beyond the scope of the book. I will, however, sketch one highly speculative model, which is a theological analogue to a Cartesian model of perception.

On the Cartesian model of perception, our sense impressions are on display in something we might call the theater of the mind. We should not suppose, however, that the mind in observing the sense impressions always forms further representations of these sense impressions, for if it always did that, we would have a vicious regress of impressions.¹⁰ Rather, the mind, without having further representations of sense impressions caused in it, directly perceives the sense impressions that are on the mental theater's stage. The sense impressions are explanatorily prior to our perceivings in a constitutive rather than causal way: our act of perceiving saltiness is constituted by an activity of the mind plus a salty sense impression.

While I am skeptical whether the Cartesian model provides a correct story about *our* perception, we can use it to model God's knowledge of the world, by taking God to be analogous to the soul and the events of the world to be the contents of the stage. (We don't need the stage itself in the model.) The events of the world then do not cause impressions or thoughts in God. Rather they are explanatorily prior to divine beliefs in a constitutive way: God's beliefs about contingent events are constituted by the events that they are about together with the activity of God's mind. This does require the rejection of an extension of causal finitism to certain kinds of constitutive explanations.

And there is independent reason in classical theism to opt for a model like this. Classical theism endorses divine simplicity. But divine simplicity implies that God has no accidental intrinsic properties. But, necessarily, God believes all and only the true propositions. Thus, the property of God believing that *p*, where *p* is a contingent truth, is a property that God will lack in worlds where *p* is not true, and hence is an accidental property. It cannot, thus, be an intrinsic property. Hence, the classical theist has to say that God's beliefs about contingent truths are partially constituted by items

¹⁰ In introspection, there may also be impressions of impressions, but these will also be on display in the theater.

external to God. And it is particularly elegant to take the facts that the beliefs are about as those items, which then gives us the above model (cf. Pruss 2008).

Hence, once again, while causal finitism yields a *prima facie* difficulty for theism, *classical* theism gives one independent motivation for a solution.

3.4 Divine action

But beyond divine knowledge and belief, there is the question of divine action. For instance, if God knows *p*, surely he can announce *p*, since he is all powerful—no one can stop him from speaking.¹¹ And now we will have a dilemma for the causal finitist.

Suppose that God knows about some infinite number of events, and he announces a fact essentially depending on that knowledge. For instance, perhaps he announces that of the infinite number of dice rolled, all but finitely many came up sixes, to recall an example from Chapter 5. Then either God's announcement is or is not caused by the infinite number of events.

If God's announcement is caused by the infinite number of events, we have a violation of causal finitism. But if God's announcement is not caused by it, then the scenario we just gave is not ruled out by causal finitism. However, such divine announcements are sufficient for many of the probability and rationality paradoxes that were used to argue for causal finitism. Hence, causal finitism does not do the job it needs to do to get us out of paradox.

In response, I bite a bullet by denying that God can always make such an announcement. To soften the blow, note first that there are questions that have an answer but where even a perfect being cannot correctly give that answer. For instance, no one can correctly answer the question:

- (9) What is an example of an entity that no speech act ever individually refers to?

But it is quite likely that there are many correct answers—surely most of the 10^{80} or more particles in the universe are never individually referred to in a speech act. There are even yes or no questions that cannot be answered correctly:

- (10) When you next answer a question, will the answer be negative? (Rescher 2005, p. 17)

If the speaker says "Yes", she answers falsely, and if she says "No", she also answers falsely. But if she refrains from answering, then as long as there is a next thing she says, there is a correct answer. For a third example, if I am so contrarian that I cannot possibly do what I've just heard predicted I will do, then no one can give me a correct answer to the question:

- (11) Am I about to clap my hands?

¹¹ I am grateful to Miguel Berasategui for raising this concern.

But the question

- (12) Did infinitely many dice show non-sixes?

differs from the above questions, in that it is possible for a being other than God to answer it. Indeed, on probabilistic grounds a sensible answer from any ordinary person who hasn't seen the dice but who knows there were infinitely many rolls is: "No". But God cannot base his answers on merely probabilistic grounds, since in doing so he would risk answering falsely, and it's impossible for a perfect being to answer falsely (cf. Kvanvig 1996).

Nonetheless, one can find other questions that it is impossible for God to answer correctly, but that a creature in principle could correctly answer. For instance:

- (13) What is an example of an entity that no divine speech act ever individually refers to?

Perhaps, for instance, God never talks about Seabiscuit with anyone, and then Seabiscuit is an example of such an entity. So there may well be a correct answer, but it is an answer God cannot give, though creatures can.

Or suppose that God promised me to keep something completely a secret from you, and the promise is fully in force. You ask God:

- (14) What did you promise Pruss to keep completely secret from me?

Being morally perfect, God cannot answer this question. Yet there is an answer, and the relevant facts wouldn't be affected by God's answering the question.

And just as there are questions a perfect being can't answer, there are true announcements that a perfect being can't make. My response disputes the claim that God could announce answers to questions like (12) even if he knew the answers, and hence sidesteps the dilemma about the causal relationship between God's announcement and the events it's about. Still a general version of question posted by the dilemma remains. When God does make announcements about some contingent facts not determined by God, is God's announcement caused by the contingent facts?

But both a positive and a negative answer can be made to fit with causal finitism. Suppose first that God's announcement is caused by the contingent facts. Then we have a simple and easy answer to why God's announcements can't be used to run paradoxes like those I have given in this book: such announcements would be just as much a violation of causal finitism as non-divine announcements or outputs from machines would be. But in any case the causal answer has a problem independent of causal finitism: considerations of aseity give one some reason to doubt that creaturely events can cause divine actions.

Next suppose that God's announcement isn't *caused* by the contingent facts. Here the causal finitist has a choice. One option is that God cannot make those announcements on infinite matters that lead to paradox, but God can make unparadoxical

announcements. While it would be *ad hoc* if we simply ruled out all the paradoxes one by one without invoking a single covering principle like causal finitism, in the case of God's involvement perhaps this is not *ad hoc*. For maybe God's perfection would not permit him to place a person in a situation where there is a rationality paradox. God is himself a rational being, rational beings are made in God's image, and to act irrationally is to act in some sense against God. Thus, perhaps, God couldn't put a person in a position where rationality required two incompatible courses of action.

But the option that appears better may be to say that the relation between the contingent matter and God's announcement is an explanatory relation analogous to a causal one. It is a relation of counterfactual dependence that is partly causally grounded, in that the worldly events involved in God's announcement (a sound booming through the air or a thought in a creature's mind) are presumably caused by God.¹² Again, here I can only sketch something that is an avenue for interesting future research.

Arguably, explanatory relations come in at least two varieties. First, there are *constitutive* or *grounding* explanations. In those cases there is a very intimate relationship between the explanans and explanandum. The explanandum here holds in virtue of the explanans. The knife is hot because its molecules have high kinetic energy—that is what constitutes it as hot, what grounds its being hot. Donald Trump is president in virtue of being validly elected. The facts reported by the explanandum and explanans are in some intuitive sense not really separate. At least in typical cases of constitutive explanations, the explanans entails the explanandum. If we take the external constitution solution to the problem of divine knowledge given in Section 3.3, we should not extend causal finitism to these kinds of explanations.

But, second, there are explanations which typically also support counterfactuals but where the facts reported by the explanans and explanandum are really separate. Typically, there is no entailment from explanans to explanandum. A paradigm case of this type of explanation is based on efficient causation. But causation may not be the only instance of this type. We could call the relations underlying this kind of explanation "quasi-causal". If we do that, then we can extend causal finitism to quasi-causal finitism, ruling out infinite quasi-causal histories, and we can say that God's announcements of contingent states of affairs when these states of affairs aren't determined by God are quasi-caused by these states of affairs. On this solution, quasi-causal finitism will be the overall view defended by the arguments of this book. This meshes nicely with the extension to causally-grounded counterfactual relations in Chapter 7, Section 2.3.

Taking the quasi-causality route, then, we can deny that God can make the paradoxical announcements, because that would violate quasi-causal finitism, and we should accept quasi-causal finitism as a natural extension of causal finitism. This seems to be the best solution to the tension between theism and the relevant paradoxes.

¹² But see Pearce (2017).

3.5 *Limits on metaphysical possibility*

The necessary existence of a perfect being imposes limits on what is possible. For instance, if such a being exists, it is impossible for there to be unredeemed evils so bad that it would be wrong for a perfect being to permit them.¹³

But such limits also limit the applicability of the kinds of rearrangement arguments that have been heavily used in this book. For instance, I argued in Chapter 3, Section 3.3 that if it is metaphysically possible to have infinitely many Grim Reapers with activation times set in some non-paradoxical way (say, at 10:00, 10:30, 10:45, and so on), then it should be possible to have the activation times set in a paradoxical way.

Now consider worlds that apart from God contain two people, Smith and Jones, and a machine, a Grim Punisher. Smith freely punches Jones in the face in a way that does not cause serious pain or injury, and never does anything else that is wrong. The Grim Punisher then imposes on Smith a certain length of moderate toothache as a punishment for this attack. The Grim Punisher has a dial that controls how long the imposed toothache will be. Clearly a world where the dial is set to some moderate length of time proportionate to Jones' injury, say a minute, is possible. But given theism, it would not appear possible to have a world where the dial is set to a billion years—at least not without introducing something else to the story, such as benefits to Smith from suffering a billion years' worth of toothache. Considerations of justice in a world with God thus limit possibilities.

Theism, thus, requires rejection of the idea that mathematically coherent modifications of possible arrangements are always going to be possible. But that idea seems important to a lot of the arguments of this book.

However, it is very plausible given theism that the plausible limits that theism places on modifications of arrangements will be based on one of two aspects of divine perfection. (For an account of omnipotence compatible with the sorts of limits I will discuss, see Pearce and Pruss 2012.) The first limit is that everything but God is created by God. This limits the possibilities for uncreated entities to one, namely God, but since none of the main arguments depend on uncreated entities, the arguments are unimpeded.

The second limit is imposed by God's moral perfection. This puts limits on examples involving suffering. But in order to be a real option, theism had better be compatible with the massive amounts of suffering found in the actual world, and hence the limits on examples involving suffering had better not be as onerous as one might at first think.

Still, there will be concerns in the case of examples in Chapter 5 that actually or potentially involve an infinite amount of pain. However, even if this makes the games incompatible with theism, they can be modified to make it possible for a perfect being to allow them. For instance, in the single-player die-guessing game, I supposed that the game has been played for an infinite amount of time and that each time you

¹³ Cf. the modal argument from evil of Gulesarian (1983).

guessed wrong you got an electric shock. On guessing strategies that do not guarantee guessing right almost always, this involves an infinite amount of total pain. And it is essential to the story that there be a real possibility of adopting such a strategy. It is, however, possible to modify the example to avoid moral worries. We might, for instance, suppose that the amount of pain caused by the shock is sufficiently small that the pleasure of the excitement of the game—regardless of whether one wins or loses—outweighs the harm.¹⁴ Granted, such a modification makes the total payoff be infinite no matter what happens. But the excitement of the game is not dependent on one's decision about how to bet, or so we may suppose, and benefits, even infinite benefits, that do not depend on the decision made should not affect the rationality of the decision—they can be bracketed off.¹⁵

Could there be other limits that theism places on possible worlds, over and beyond those coming from God's being creator and perfectly morally good? Maybe, but it is plausible that there cannot be too many limits since among God's perfections there is, after all, his infinite power which should leave in place many options (Pruss 2016).

In any case, we would do well to remember that causal and quasi-causal finitism themselves place a limitation on rearrangements. As noted in Chapter 1, Section 3.2, what we want to avoid are *ad hoc* limitations on rearrangements. The theistic limitations, just as the limitation due to causal finitism, are principled and not *ad hoc*.

4. Evaluation

Causal finitism implies the existence of at least one first cause. Given a Causal Principle for contingent beings, it follows that any first cause is a necessary being. To argue from a necessary first cause to theism would be a major task, a task attempted by natural theologians like Aquinas (1920) and Clarke (1803), but beyond the scope of this book.

Nonetheless, theism is the best developed extant theory of a necessary first cause. Therefore, if theism were incompatible with causal finitism, that would make causal finitism less plausible. I have argued that theism can be reconciled with causal finitism, especially if theism is understood in a classical way. We may, however, have to extend causal finitism to quasi-causal finitism.

Perhaps the main alternative to theism here would be the claim that the Big Bang event is a necessarily obtaining first cause. Defending this claim would force one to deny the intuition that the laws of nature and initial arrangement of matter could have been very different from how they in fact are.

¹⁴ Or we might suppose that the player each time freely chooses to play in order to save a friend from harm. God can then be justified in allowing the game to go on because the player grows in virtue through playing it, despite the pain suffered.

¹⁵ If such bracketing cannot be done, then a strongly egalitarian universalist theology on which everyone receives an equal infinite bliss in the afterlife would undercut all ordinary decisions by ensuring the same payoff no matter what, but it does not seem to do so (one might use hyperreal utilities to explain why not—cf. Herzberg 2011).

Conclusions

A large number of paradoxes of infinity has been surveyed. The paradoxes differ widely. Some involve deterministic processes and others indeterministic ones. One (Thomson's Lamp) is paradoxical as it threatens violation of the Principle of Sufficient Reason. Some others threaten contradiction. Yet others threaten violation of principles of rationality, some epistemic and some practical. But there is one thing all the paradoxes have in common: an infinite number of items are in some sense causally prior to a single item.

One can always get out of any one of the individual paradoxes simply by saying that the situation it describes is impossible *because* it leads to paradox. But this is not satisfactory. First, the move tends to violate plausible instances of rearrangement, because there are similar setups that do not involve paradox. Second, the move ends up giving a wide variety of explanations for why the setups are impossible, and a unified explanation would be preferable. Causal finitism provides such a unified explanation by denying the possibility of an infinite causal history for a single item.

Causal finitism walks a line between finitism *simpliciter* and what one might call unrestricted infinitism, on which the finite/infinite division does not by itself mark a difference in possibility, whether in causal or non-causal situations. Finitism, when combined with either the growing block theory of time or eternalism, can resolve all our paradoxes, but the cost of that resolution is high. First, finitism implies that modern mathematics concerns itself with impossible situations. Second, finitism plus eternalism contradicts the plausible claim that future infinities are possible. And while finitism plus growing block avoids the paradoxes that motivate causal finitism and allows for an infinite future, it leads to skepticism about what time is now. On the other hand, unrestricted infinitism does not allow for a unified solution to the paradoxes.

In addition to solving paradoxes, we have seen that causal finitism does justice to common intuitions about the viciousness of infinite causal regresses. And, as a bonus, it gives us a potentially interesting metaphysical account of the finite and the countable with an application to the philosophy of mathematics.

Causal finitism should be thought of as a family of views, depending on what causal priority relations are used to define it. Further research is called for to specify causal finitism further, as well as to consider plausible analogues or extensions to causal finitism. For instance, one might argue that just as nothing can have infinitely many causes, nothing can have infinitely many parts. There are also intriguing

analogues between infinite causal histories and circular causal histories, and causal finitism can be extended to preclude both.

There is an initially plausible argument from causal finitism to the thesis that time (and probably then also space) are discrete. However, there are empirically adequate interpretations of causality in physics on which causal sequences are discrete but time is not. If these interpretations are implausible, then causal finitism does push towards temporal discretism.

On the other hand, it definitely does follow from causal finitism together with the denial of the possibility of causal loops that there must be an uncaused cause. Investigation is needed whether this uncaused cause is something natural, like the Big Bang, or something supernatural, like God. At the same time, while the arguments for causal finitism have a similarity to parts of Kalām arguments for the existence of God, there are special difficulties—ones that I believe can be overcome—that have to be faced for someone who combines causal finitism with theism. Pruss and Rasmussen (2018) have argued that theists and non-theists could agree that there is a necessary being and investigate the nature of this being together. The Kalām-inspired argument of Chapter 9 should help further motivate this joint investigation.

References

- Albert, David Z. 1992. *Quantum Mechanics and Experience*. Cambridge, MA: Harvard University Press.
- Albert, David Z. and B. Loewer. 1988. "Interpreting the Many Worlds Interpretation". In *Synthese* 77: 195–213.
- Alexander, David E. and Daniel M. Johnson, eds. 2016. *Calvinism and the Problem of Evil*. Eugene, OR: Wipf and Stock.
- Anscombe, G. E. M. 1971. *Causality and Determination: An Inaugural Lecture*. Cambridge: Cambridge University Press.
- Aquinas, Saint Thomas. 1920. *Summa Theologica*, trans. by Fathers of the English Dominican Province. London.
- Aristotle. 1934. *Physics*: Books 5–8, trans. by P. H. Wicksteed and F. M. Cornford. Cambridge, MA: Harvard University Press.
- Aristotle. 1957. *Physics*: Books 1–4, trans. by P. H. Wicksteed and F. M. Cornford. Cambridge, MA: Harvard University Press.
- Aristotle. 1984. *Complete Works*, ed. by Jonathan Barnes. Princeton, NJ: Princeton University Press.
- Arntzenius, F., A. Elga, and J. Hawthorne. 2004. "Bayesianism, Infinite Decisions, and Binding". In *Mind* 133: 251–83.
- Artin, Michael, Alexandre Grothendieck, and Jean-Louis Verdier, eds. 1972. "Seminaire de Geometrie Algebrique du Bois Marie, 1963–64: Theorie des topos et cohomologie etale des schemas (SGA 4)". *Lecture Notes in Mathematics* 269. Berlin: Springer.
- Barrett, Jeffrey A. 1995. "The Single-Mind and Many-Minds Versions of Quantum Mechanics". In *Erkenntnis* 42: 89–105.
- Bartha, Paul. 2011. "Symmetry and the Brown–Freiling Refutation of the Continuum Hypothesis". In *Symmetry* 3: 636–52.
- Bartha, Paul and Richard Johns. 2001. "Probability and Symmetry". In *Proceedings of the Philosophy of Science Association*: S109–S122.
- Beam, J. 2007. "Unfair Gambles in Probability". In *Statistics and Probability Letters* 77: 681–6.
- Bell, John S. 1987. "Are There Quantum Jumps?" In *Schrodinger: Centenary Celebration of a Polymath*, ed. by C. W. Kilmister. Cambridge: Cambridge University Press, pp. 41–52.
- Benacerraf, Paul. 1962. "Tasks, Super-Tasks, and the Modern Eleatics". In *Journal of Philosophy* 59: 765–84.
- Benacerraf, Paul. 1965. "What Numbers Could Not Be". In *Philosophical Review* 74: 47–73.
- Benardete, J. 1964. *Infinity: An Essay in Metaphysics*. Oxford: Clarendon Press.
- Bohm, D. 1952. "A Suggested Interpretation of the Quantum Theory in Terms of 'Hidden' Variables, I and II". In *Physical Review* 85: 166–93.
- Boolos, George S. and Richard C. Jeffrey. 1995. *Computability and Logic*, 3rd edn. Cambridge: Cambridge University Press.
- Bradley, F. H. 1893. *Appearance and Reality: A Metaphysical Essay*. London: Swann Sonnenschein.

- Buchak, Lara. 2014. *Risk and Rationality*. Oxford: Oxford University Press.
- Carrier, Martin. 2003. "How to Tell Causes from Effects: Kant's Causal Theory of Time and Modern Approaches". In *Studies in History and Philosophy of Science, Part A* 34: 59–71.
- Chalmers, David. 2002. "The St. Petersburg Two-Envelope Paradox". In *Analysis* 62: 155–7.
- Clarke, Samuel. 1803. *A Discourse Concerning the Being and Attributes of God, the Obligations of Natural Religion, and the Truth and Certainty of the Christian Revelation*. Glasgow: Griffin.
- Craig, William Lane. 1979. *The Kalām Cosmological Argument*. Cambridge: Cambridge University Press.
- Craig, William Lane. 2009. "The Kalam Cosmological Argument". In *The Blackwell Companion to Natural Theology*, ed. by W. L. Craig and J. P. Moreland. Malden, MA: Wiley-Blackwell, pp. 101–201.
- Davidson, Donald. 1963. "Actions, Reasons, and Causes". In *Journal of Philosophy* 60: 685–700.
- Della Rocca, Michael. 2010. "PSR". In *Philosophers' Imprint* 10: 1–13.
- Dewdney, Alexander. 2001. *The Planiverse: Computer Contact With a Two-Dimensional World*. New York: Copernicus.
- Diekmeyer, Joseph. 2014. "The Existence of the Past". In *Synthese* 191: 1085–104.
- Dirac, P. A. M. 1987. *The Principles of Quantum Mechanics*, 4th edn. Oxford: Clarendon Press.
- Dougherty, M. V. 2011. *Moral Dilemmas in Medieval Thought: From Gratian to Aquinas*. Cambridge: Cambridge University Press.
- Dougherty, Trent and Alexander R. Pruss. 2014. "Evil and the Problem of Anomaly". In *Oxford Studies in the Philosophy of Religion*, Vol. 5, ed. by J. Kvanvig. Oxford: Oxford University Press, pp. 49–87.
- Dubins, L. E. 1975. "Finitely Additive Conditional Probabilities, Conglomerability, and Disintegration". In *Annals of Probability* 3: 89–99.
- Dummett, Michael. 1986. "Causal Loops". In *The Nature of Time*, ed. by Raymond Flood and Michael Lockwood. Oxford: Blackwell, pp. 135–69.
- Easwaran, Kenny. 2013. "Why Countable Additivity?" In *Thought* 2: 53–61.
- Edwards, Paul. 1959. "The Cosmological Argument". In *The Rationalist Annual for 1959*. London: Pemberton, pp. 63–77.
- Einstein, A. 2015. *Relativity: The Special and the General Theory*, trans. by Robert W. Lawson. Princeton: Princeton University Press.
- Elga, Adam. 2000. "Self-Locating Belief and the Sleeping Beauty Problem". In *Analysis* 60: 143–7.
- Everett, H. 1957. "'Relative State' Formulation of Quantum Mechanics". In *Reviews of Modern Physics* 29: 454–62.
- Feynman, Richard. 1985. *Surely You're Joking, Mr. Feynman!* New York: Norton.
- Fine, Terence L. 1973. *Theories of Probability: An Examination of Foundations*. New York: Academic Press.
- Fitelson, Branden and Alan Hájek. 2017. "Declarations of Independence". In *Synthese* 194: 3979–95.
- Freiling, C. 1986. "Axioms of Symmetry: Throwing Darts at the Real Number Line". In *Journal of Symbolic Logic* 51: 190–200.
- Ghirardi, G. C., A. Rimini, and T. Weber. 1986. "Unified Dynamics for Microscopic and Macroscopic Systems". In *Physical Review D* 34: 470.

- Good, I. J. 1967. "On the Principle of Total Evidence". In *British Journal for the Philosophy of Science* 17: 319–21.
- Greaves, H. 2006. "Probability in the Everett Interpretation". In *Philosophy Compass* 2: 109–28.
- Gulesarian, Theodore. 1983. "God and Possible Worlds: The Modal Problem of Evil". In *Noûs* 17: 221–38.
- Hagar, Amit. 2014. *Discrete or Continuous? The Quest for Fundamental Length in Modern Physics*. Cambridge: Cambridge University Press.
- Haldane, John. 2007. "Privative Causality". In *Analysis* 67: 180–6.
- Hansen, Casper Storm. n.d. "A Fair Infinite Lottery". <https://pdfs.semanticscholar.org/9184/92ab5d678ba22a10270b4a500f286a99505b.pdf>.
- Hardin, Christopher S. and Alan T. Taylor. 2008. "An Introduction to Infinite Hat Problems". In *The Mathematical Intelligencer* 30: 20–5.
- Hawthorne, John. 2000. "Before-Effect and Zeno Causality". In *Noûs* 34: 622–33.
- Hellman, Geoffrey. 1981. "How to Gödel a Frege-Russell: Gödel's Incompleteness Theorems and Logicism". In *Noûs* 15: 451–68.
- Herzberg, Frederik. 2011. "Hyperreal Expected Utilities and Pascal's Wager". In *Logique et Analyse* 54: 69–108.
- Hesslow, Germund. 1988. "The Problem of Causal Selection". In *Contemporary Science and Natural Explanation*, ed. by D. J. Hilton. New York: New York University Press, pp. 11–32.
- Hicks, Michael Townsen and Peter van Elswyk. 2015. "Humean Laws and Circular Explanation". In *Philosophical Studies* 172: 433–43.
- Hofmann-Wellenhoff, Bernhard and Helmut Moritz. 2006. *Physical Geodesy*, 2nd edn. Vienna: Springer.
- Howson, Colin. 2014. "Finite Additivity, Another Lottery Paradox and Conditionalisation". In *Synthese* 191: 989–1012.
- Huemer, Michael. 2016. *Approaching Infinity*. New York: Palgrave Macmillan.
- Hume, David. 1779. *Dialogues Concerning Natural Religion*. Cambridge: Cambridge University Press.
- Jackson, Allyn. 2008. "Interview with Martin Davis". In *Notices of the American Mathematical Society* 55: 560–71.
- Jech, T. J. 1973. *The Axiom of Choice*. Amsterdam: North Holland.
- Kadane, J. B., M. J. Schervish, and T. Seidenfeld. 1996. "Reasoning to a Foregone Conclusion". In *Journal of the American Statistical Association* 91: 1228–35.
- Kane, Robert. 1999. *The Significance of Free Will*. Oxford: Oxford University Press.
- Kant, Immanuel. 1907. *Immanuel Kant's Critique of Pure Reason*, trans. by F. Max Müller. London: Macmillan.
- Klein, P. 1998. "Foundationalism and the Infinite Regress of Reasons". In *Philosophy and Phenomenological Research* 58: 919–25.
- Knuth, Donald E. 1974. *Surreal Numbers: How Two Ex-Students Turned on to Pure Mathematics and Found Total Happiness: A Mathematical Novelette*. Reading, MA: Addison-Wesley.
- Knuth, Donald E. 2002. "All Questions Answered". In *Notices of the American Mathematical Society* 49: 318–24.
- Koons, Robert. 2014. "A New Kalam Argument: Revenge of the Grim Reaper". In *Noûs* 48: 256–67.

- Krauss, P. H. 1968. "Representation of Conditional Probability Measures on Boolean Algebras". In *Acta Mathematica Academiae Scientiarum Hungarica* 19: 229–41.
- Kunen, Kenneth. 1980. *Set Theory: An Introduction to Independence Proofs*. New York: Elsevier.
- Kvanvig, Jonathan. 1996. *The Possibility of an All-Knowing God*. New York: Palgrave Macmillan.
- Lang, Serge. 2002. *Algebra*, 3rd edn. Dordrecht: Springer.
- Leibniz, G. W. 1989. *Philosophical Essays*, trans. by Roger Ariew and Daniel Garber. Indianapolis, IN: Hackett.
- Lewis, David. 1973. "Causation". In *Journal of Philosophy* 70: 556–67.
- Lewis, David. 1976. "The Paradoxes of Time Travel". In *American Philosophical Quarterly* 13: 145–52.
- Lewis, David. 1979. "Counterfactual Dependence and Time's Arrow". In *Noûs* 13: 455–76.
- Lewis, David. 1986. *On the Plurality of Worlds*. Oxford: Blackwell.
- Lewis, David. 2004. "Tensed Quantifiers". In *Oxford Studies in Metaphysics*, Vol. 1, ed. by Dean W. Zimmerman. Oxford: Oxford University Press, pp. 3–14.
- McDaniel, Kris. 2009. "Extended Simples and Qualitative Heterogeneity". In *Philosophical Quarterly* 59: 325–31.
- McGee, Vann. 1994. "Learning the Impossible". In *Probability and Conditionals*, ed. by Ellery Eells and Brian Skyrms. Cambridge: Cambridge University Press, pp. 179–99.
- Mellor, D. H. 1998. *Real Time II*. Abingdon: Routledge.
- Mendelson, Elliott. 2008. *Number Systems and the Foundations of Analysis*. Mineola, NY: Dover.
- Merricks, Trenton. 2001. *Objects and Persons*. Oxford: Oxford University Press.
- Merricks, Trenton. 2006. "Good-Bye Growing Block". In *Oxford Studies in Metaphysics* 2: 103–10.
- Meyer, Ulrich. 2012. "Explaining Causal Loops". In *Analysis* 72: 259–64.
- Nielsen, Lawrence E. and Robert F. Landel. 1994. *Mechanical Properties of Polymers and Composites*, 2nd edn. New York: Dekker.
- Norton, John D. 2018. "How to Build an Infinite Lottery Machine". In *European Journal for Philosophy of Science* 8(1): 71–95.
- Norton, John D. and Alexander R. Pruss. 2018. "Correction to John D. Norton 'How to Build an Infinite Lottery Machine'". In *European Journal for Philosophy of Science* 8: 143–4.
- Oppy, Graham. 2006. *Philosophical Perspectives on Infinity*. Cambridge: Cambridge University Press.
- Parsons, J. 2000. "Must a Four-Dimensionalist Believe in Temporal Parts?" In *The Monist* 83: 399–418.
- Pascal, Blaise. 1858. *Oeuvres Complètes*, Vol. 1. Paris: Hachette.
- Pearce, Kenneth L. 2017. "Counterpossible Dependence and the Efficacy of the Divine Will". In *Faith and Philosophy* 34: 3–16.
- Pearce, Kenneth L. and Alexander R. Pruss. 2012. "Understanding Omnipotence". In *Religious Studies* 48: 403–14.
- Plato. 1996. *Parmenides*, trans. by Mary Louise Gill and Paul Ryan. Indianapolis, IN: Hackett.
- Pruss, Alexander R. 1998. "The Hume–Edwards Principle and the Cosmological Argument". In *International Journal for Philosophy of Religion* 43: 149–65.
- Pruss, Alexander R. 2006. *The Principle of Sufficient Reason: A Reassessment*. Cambridge: Cambridge University Press.
- Pruss, Alexander R. 2007. "Prophecy Without Middle Knowledge". In *Faith and Philosophy* 24: 433–57.

- Pruss, Alexander R. 2008. "On Two Problems of Divine Simplicity". In *Oxford Studies in the Philosophy of Religion*, Vol. 1, ed. by Jonathan L. Kvanvig. Oxford: Oxford University Press, pp. 150–67.
- Pruss, Alexander R. 2011. *Actuality, Possibility and Worlds*. New York: Continuum.
- Pruss, Alexander R. 2012. "The Leibnizian Cosmological Argument". In *The Blackwell Companion to Natural Theology*, ed. by W. L. Craig and J. P. Moreland. Oxford: Blackwell, pp. 24–100.
- Pruss, Alexander R. 2013a. "Null Probability, Dominance and Rotation". In *Analysis* 73: 682–5.
- Pruss, Alexander R. 2013b. "Omnirationality". In *Res Philosophica* 90: 1–21.
- Pruss, Alexander R. 2014. "Infinitesimals Are Too Small for Countably Infinite Fair Lotteries". In *Synthese* 191: 1051–7.
- Pruss, Alexander R. 2015. "Popper Functions, Uniform Distributions and Infinite Sequences of Heads". In *Journal of Philosophical Logic* 44: 259–71.
- Pruss, Alexander R. 2016. "Divine Creative Freedom". In: *Oxford Studies in the Philosophy of Religion*, Vol. 7, ed. by Jonathan L. Kvanvig. Oxford: Oxford University Press, pp. 213–38.
- Pruss, Alexander R. 2017. "The Principle of Sufficient Reason and Probability". In *Oxford Studies in Metaphysics*, Vol. 10, ed. by Karen Bennett and Dean Zimmerman. Oxford: Oxford University Press, pp. 261–78.
- Pruss, Alexander R. 2018. "A Traveling Forms Interpretation of Quantum Mechanics". In: *Neo-Aristotelian Perspectives on Contemporary Science*, ed. by W. M. R. Simpson, R. C. Koons, and N. J. Teh. New York: Routledge, pp. 105–22.
- Pruss, Alexander R. ms. "Might All Infinities Be The Same Size?"
- Pruss, Alexander R. and Joshua Rasmussen. 2018. *Necessary Existence*. Oxford: Oxford University Press.
- Rescher, Nicholas. 2005. *Studies in Philosophical Inquiry*. Frankfurt: Ontos Verlag.
- Rescher, Nicholas. 2013. *G. W. Leibniz's Monadology: An Edition for Students*. Abingdon: Routledge.
- Robinson, Abraham. 1996. *Non-Standard Analysis*, revised edn. Princeton, NJ: Princeton University Press.
- Rowe, W. L. 1970. "Two Criticisms of the Cosmological Argument". In *The Monist* 54: 441–59.
- Rudin, Walter. 1987. *Real and Complex Analysis*, 3rd edn. New York: McGraw-Hill.
- Schaffer, Jonathan. 2016. "Grounding in the Image of Causation". In *Philosophical Studies* 173: 49–100.
- Schervish, M. J., T. Seidenfeld, and J. B. Kadane. 1984. "The Extent of Non-Conglomerability of Finitely Additive Probabilities". In *Probability Theory and Related Fields* 66: 205–26.
- Schmid, Christoph. 2012. "Newton's Superb Theorem: An Elementary Geometric Proof". <https://arxiv.org/abs/1201.6534>.
- Shackel, Nicholas. 2005. "The Form of the Benardete Dichotomy". In *British Journal for the Philosophy of Science* 56: 397–417.
- Shapiro, Stewart. 1985. "Second-Order Languages and Mathematical Practice". In *Journal of Symbolic Logic* 50: 714–42.
- Sheehan, David. 2009. *Physical Biochemistry: Principles and Applications*, 2nd edn. Chichester: Wiley.

- Silva, Tomás Oliveira e. 2015. "Goldbach Conjecture Verification". <http://sweet.ua.pt/tos/goldbach.html>.
- Smith, Quentin. 1993. *The Language of Time*. New York: Oxford University Press.
- Smullyan, R. 2008. "Memories and Inconsistencies". In *A Lifetime of Puzzles: A Collection of Puzzles in Honor of Martin Gardner's 90th Birthday*, ed. by E. D. Demaine, M. L. Demaine, and T. Rodgers. Wellesley, MA: A. K. Peters, pp. 245–52.
- Solovay, R. M. 1970. "A Model of Set-Theory in Which Every Set of Reals is Lebesgue Measurable". In *Transactions of the American Mathematical Society* 92: 1–56.
- Taliaferro, Charles. 1993. "Unknowable Truths and Omniscience: A Reply to Kvanvig". In *Journal of the American Academy of Religion* 61: 553–66.
- Thomson, James F. 1954. "Tasks and Super-Tasks". In *Analysis* 15: 1–13.
- Thorp, B. L. D. 1967. "Solution of Problem 5348". In *American Mathematical Monthly* 72: 730–1.
- van Fraassen, Bas C. 1984. "Belief and the Will". In *Journal of Philosophy* 81: 235–56.
- van Inwagen, Peter. 1995. *Material Beings*. Ithaca, NY: Cornell University Press.
- Wagon, Stan. 1994. *The Banach–Tarski Paradox*. Cambridge: Cambridge University Press.
- Wapner, Leonard M. 2007. *The Pea and the Sun: A Mathematical Paradox*. London: Taylor & Francis.
- Weatherson, Brian and Dan Marshall. 2014. "Intrinsic vs. Extrinsic Properties", *The Stanford Encyclopedia of Philosophy*, ed. by Edward N. Zalta. <https://plato.stanford.edu/archives/fall2014/entries/intrinsic-extrinsic/>.
- Weirich, Paul. 1984. "The St. Petersburg Gamble and Risk". In *Theory and Decision* 17: 193–202.
- Williamson, Timothy. 2007. "How Probable Is an Infinite Sequence of Heads?" In *Analysis* 67: 173–80.
- Wiseman, Richard. 2002. "LaughLab: The Scientific Search for the World's Funniest Joke". <https://richardwiseman.files.wordpress.com/2011/09/ll-final-report.pdf>.
- Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.

Index

- abstracta 4, 6, 13–14, 185
- ACCR 119–25
- action 99, 105, 108, 143, 164–5, 170n2, 171,
 - 176, 179, 184–7, 190
 - divine 185, 188–9
 - group 135, 138
 - instantaneous 47n8, 48, 177
 - moral 144
- angels 98–104, 106, 130–1, 135–6, 143
- Anscombe, G. E. M. 172
- Aquinas, Saint Thomas 26, 106, 192
- Archimedean property of real numbers 41, 64n1
- Aristotle 10, 87, 164, 168–9, 178, 180. *See also* metaphysics, Aristotelian
- arithmetic 16, 20–2, 41, 65, 164
 - axioms of 16, 20. *See also* Peano Axioms
 - Peano 22
 - second-order 18n16
- atemporality 152, 170
- Axiom of Choice 3, 14, 18n17, 25, 38, 41n2, 65n2, 83–4, 88–9, 97–8, 101, 117–39, 183n6
 - Countable 5n3
 - for countable collections of reals 119–25
 - paradoxes of 117–39
- Axiom of Countable Additivity 81
- Axiom of Dependent Choice 38–9, 70n5, 136, 182n1, 183n5. *See also* Axiom of Choice
- Axiom of Finitude 14
- Axiom of Regularity 36, 162n8
- Axiom of Separation 24n23
- axiomatization 21
- axioms 14–18, 20–2, 121
 - of set theory 22, 91, 117–18. *See also* Zermelo–Fraenkel axioms
- Benacerraf, Paul 1–2, 20, 43–4
- Benacerraf’s Identification Problem 20, 22
- Benardete, José 146. *See also* paradox,
 - Benardete’s Boards
- Big Bang. *See* events, Big Bang event
- black holes 155, 157, 180
- Boolean Prime Ideal axiom 41n2
- Cantor, Georg 5. *See also* theorem, Cantor’s cardinality. *See* set, cardinality
- cardinals. *See also* set, cardinality
 - inaccessible 136
- causation 12, 27–8, 32–3, 36, 54, 142–3, 145, 149, 152, 161, 165–6, 172–4, 176–7, 180–2, 187, 190
 - Aristotelian 8, 44, 46, 148, 164, 170, 180
 - backwards 3, 63, 152, 159, 170–1
 - by absences 149, 166
 - causal finitism 1–4, 7–11, 13, 18–19, 22–3, 25, 27, 35, 37–42, 46–50, 54–7, 59–61, 63, 68–70, 81, 86–90, 92–6, 100, 103, 107–8, 110–11, 113–15, 118, 124, 135, 137n18, 140–52, 157, 159–67, 169–70, 172–4, 176–82, 184–90, 192–4
 - causal infinitism 4, 26, 48, 57, 60–1, 64, 68, 79, 81, 87n23, 88–90, 93–6, 98, 100, 103–4, 118–19, 130–2, 134–8, 160, 177
 - causal loops 31–2, 40, 61–3, 140, 162–3, 166, 181, 194
 - causal nexus 37, 50, 141, 145, 162–3
 - causal powers 7n5, 8–9, 14, 44–5, 51, 61n20, 63, 144, 147, 150, 192
 - Causal Principle 52–4, 182–4, 192
 - causal priority 142–3, 146–7, 163, 165, 193
 - causal regresses 25–7, 29, 31, 33, 35–7, 151, 162, 193
 - causal relations 35, 37, 55, 140, 141–3, 145–9, 162, 165–6, 171, 189, 193
 - counterfactual theory of 51, 54
 - derivatives 32–3, 142
 - deterministic 51
 - discrete 174, 177–80
 - efficient 30, 164, 190
 - final 164–5
 - first cause 37, 181–92. *See also* God, theism
 - formal 164
 - fundamental 32–3, 37, 141–2, 160, 162, 165
 - Humean 9
 - indeterministic 34, 51, 177
 - manipulationist theories of 54
 - material 148, 164
 - microphysical 142
 - partial 37, 142–5, 147–8, 162–3, 165–6
 - quantum 176–7, 179. *See also* physics,
 - Quantum Mechanics
 - quasi-causal 181, 190, 192
 - self-causation 23, 36n9, 148, 163
 - simultaneous 151, 159, 170
- Chalmers, David 68–9
- change 43–4, 168
 - accidental 170
 - instant 155

- change (*Cont.*)
 time as the measure of 168, 171
 classical mechanics 167
 classical thermodynamics 176
 compatibilism
 theological 186
 concreta 4, 13–14
 conglomerability 75–7, 79, 101n4
 contingency 22n22, 43, 52, 56, 91, 122,
 124, 152, 164n10, 181–3, 186–7,
 189–90, 192
 continuous distributions 87–8, 92
 continuum 87, 101, 119, 125, 132, 135. *See also*
 set, uncountable
 continuum hypothesis 67n3
 countable additivity 75–7, 81, 87n24, 96n2,
 121, 138
 Craig, William Lane 15, 151, 184
 credences 70–2, 74–6, 100–1. *See also*
 probability, epistemic
 posterior 75

 decision theory 3, 93–116, 137
 density 57–8, 60, 153, 155, 157–8
 gradient 155
 infinite 153, 155
 of time 29
 uniform 57n14, 59, 155
 desire satisfaction 110
 determinism 48n10, 174
 dispositions 136, 150
 dominance arguments 68–9, 107
 Dummett, Michael 30
 Dutch Books 68–9, 70–1, 118, 128–30,
 134, 138
 infinite 66

 Einstein, Albert 131
 emergence
 necessary 52
 equations 75–6
 differential 91, 172
 partial differential 178n6
 Schrödinger 174–5, 177–8. *See also* physics,
 Quantum Mechanics
 second-order 178n6
 equivalence classes 84, 88–9, 98, 120,
 134–5, 139
 essentiality of origins 8, 35, 144
 essentiality of time of origination
 104
 eternalism. *See* time, eternalist theory of
 eternity 6, 164
 atemporality 152, 170
 eternal life 40, 60–1, 63
 omnitemporality. *See* omnitemporality
 timelessness 54

 events 9–10, 26, 30n6, 40, 46, 51, 61, 65–6,
 76–7, 81, 83, 96n2, 131, 145, 149–50, 161,
 164, 170, 176, 186, 187, 189–90
 Big Bang event 180, 182, 192, 194
 contingent 52, 181, 186–7
 disjoint 75–6
 disjunctive 28
 fine-grained 141–2, 165
 future 10, 17–18, 23, 49–50, 103, 150
 individuation of 140–2
 infinitely many 49–50, 108, 131, 188
 mental 160
 past 23, 63
 physical 89, 183
 positive 146, 149
 quantum 54n11
 random 122, 135
 stochastic 187
 uncaused 51–3, 54n11, 184
 zero probability 124
 explanation 1–2, 4, 27, 29–30, 37, 43–4, 100,
 108, 154, 156, 161–2, 164–6, 173–5, 190,
 193
 causal 29, 34–5, 173
 constitutive 187, 190
 economic 154
 explanatory circularity (loops) 29–32
 explanatory priority 30, 187
 explanatory relations 164–5, 187, 190
 external 29
 grounding 30n6, 190
 infinite 140
 internal 30
 metaphysical 30n6
 nomic 30n6
 of regresses 29
 of the universe 26, 37
 partial 164–5
 psychological 154
 scientific 30n6, 153, 155–6, 193
 self- 29
 unified 161, 165

 Fermat's Last Theorem 20
 Feynman, Richard 125
 fields 135–6, 161, 178–80
 electromagnetic 135, 178–9
 gravitational 53, 131, 146, 161, 169
 magnetic 174
 physical 53
 finitism 3, 10–19, 22–3, 48–50, 61, 103, 118–19,
 140, 148, 150, 153–4, 164, 170, 184, 193
 causal. *See* causation, causal finitism
 explanatory 165
 growing block 11
 mereological 164
 presentist 10

- free will
 - indeterministic 48, 62
 - quasi- 48n10
- functions 5, 26, 81, 83–4, 87n24, 88, 95, 102, 117–18, 120, 122, 124, 127n9, 175, 179
 - choice 88, 98, 102, 117–19, 122, 124–5, 130–2, 137, 139n19
 - constructible 118
 - measurable 137
 - non-measurable 137
 - one-to-one 119, 120n3, 124, 132–5
 - Popper 82
 - probability 75–6, 85–6
 - stochastic 178
 - wavefunction 34, 174–8. *See also* physics, Quantum Mechanics
- fusion. *See* mereology, mereological sum
- geometry 168–9
 - internal 169
- God 8, 26, 33–5, 36n9, 46n7, 54, 70n6, 152, 165, 181–2, 184–92, 194. *See also* theism
 - aseity of 36n9, 187, 189
 - omniscience of 181, 186
 - perfection of 190–2
- Gödel, Kurt 16, 20. *See also* theorem, Gödel's incompleteness theorems
- Goldbach's conjecture 16–17
- graphs 25, 37, 162
 - directed 38
- grounding 9, 30n6, 32–3, 36–7, 45, 107, 142, 164, 171, 174, 179, 190
 - broadly logical 145
 - causal 145, 165, 190
 - self- 36
- group 135
 - action 138
 - countable 135, 138
 - free 138
 - permutation 138
- growing block theory of time. *See* time, growing block theory of
- Hilbert's Hotel 12, 14, 156–7
- Howson's rule 74–5
- Huemer, Michael 1, 140, 153–9, 180
- Hume, David 26, 29, 37, 48n10, 52, 55. *See also* laws of nature, Humean theory of
- Hume–Edwards Principle (HEP) 29
- if–thenism 15–17. *See also* mathematics
- infinetism. *See* finitism
- infinity 1–2, 12, 17, 23, 27, 35, 41, 47, 50, 58–9, 63, 65, 86, 90–1, 107–8, 134, 142, 144, 150–5, 158–60, 163, 165–6, 170–1, 185, 193
 - actual 13
 - causal 160, 177. *See also* causation, causal finitism
 - completed 151–2
 - countable. *See* set, countable
 - future 17–18, 113, 150–1, 170, 187, 193
 - intensive 154–6
 - past 140, 150–2, 186
 - potential 10, 15
 - simultaneous 160
 - uncountable. *See* set, uncountable
- Instant Grim Reapers (IGRs) 183–4. *See also* paradox, Grim Reaper
- integrals 89, 156
- justification
 - epistemic 27–8
- Kalām cosmological argument 50, 181, 184, 194
- Kant, Immanuel 146, 151
- knowledge 73, 88, 105n8, 109, 134–5, 151
 - divine 181, 186–8, 190
- Knuth, Donald 17, 65
- König's Lemma 39n11
- Law of Excluded Middle 171n3
- Law of Large Numbers 80, 170n2
- laws of nature 44–6, 55, 88, 131, 152, 156, 158, 172, 180, 192
 - Aristotelian theory of 8, 44–6
 - causal 9
 - Humean theory of 30
 - infinitely complex 137
 - Newtonian 60, 178n6. *See also* physics, Newtonian
 - relativistic 42n3
- Leibniz, Gottfried 22n22, 156, 185n7
- Lewis, David 7–8, 24, 51, 54, 61n20, 146n1
- logic 3–7, 14, 16, 17n14, 45, 53, 67, 100–1, 105, 145, 174
 - first-order 53
 - logically revisionary answers to paradox 1–2
 - non-classical 1
- logicism 16
- lotteries
 - paradoxical 64–92
- machines 189
 - Axiom of Choice Machine 117–39
 - constructible 81
 - four-dimensional 131–6
 - Grim Punisher 191
 - Grim Reaper. *See* paradox, Grim Reaper
 - infinite 81
 - lottery 80–1. *See also* paradox, lottery
 - three-dimensional 132, 136
 - time 62. *See also* time, travel

- magnitudes
 - extensive 153–5
 - fundamental 154
 - infinite intensive 140, 153–9
 - natural 153–6, 158
- mathematics 4–7, 13–18, 20–3, 41, 60–1, 65, 75, 77, 88–91, 94–5, 98, 111–12, 118, 121n6, 122n8, 123, 125–9, 136n16, 137, 150, 170–1, 175, 178–9, 191, 193. *See also* truth, mathematical
- infinitary 91
- mathematical definition 18
- mathematical fiction 178–9
- mathematical idealization 91
- mathematical induction 38, 115, 117n1
- mathematical limit 42, 87n24, 90, 115, 158
- mathematical object 14, 20, 122–3
- mathematical practice 14–15
- meta-mathematics 16
- philosophy of 3, 14–15, 20, 23, 63, 193
- mereology
 - infinite mereological sum 148, 164–6
 - mereological finitism 164
 - mereological sum 6, 34, 50–2, 147–8
 - mereological universalism 52
- Merricks, Trenton 11, 51, 150
- metaphysics 8, 52, 91, 150, 180
 - Aristotelian 8, 32n8, 44–6, 59, 124, 148, 151–2, 168–71, 174, 176n5, 178, 180
 - of material objects 8
 - of sets 124
 - Platonic 14–15, 23, 124, 161
- minds 19–20, 107, 176–7, 187, 190
 - angel's mind 143
 - embodied 160
 - God's mind 187. *See also* God
 - immaterial 156, 158–9, 176
 - many-minds interpretation of Quantum Mechanics 176
- modality 7–10
 - causal powers account of 7n5
 - theory of 7
- multiverse 50, 122, 130, 176
 - branching 175–7
 - deterministic 176
 - dualist 177
 - finite 177
 - infinite 118, 148, 177
 - quantum 177
- necessary being 4, 8, 35, 181–4, 192, 194. *See also* God, theism
- necessity
 - logical 16
 - metaphysical 7, 16, 157
 - narrowly logical 16, 22n22
- Newton, Isaac 4, 57n14, 172, 178n6. *See also* physics, Newtonian
- numbers 1, 10–11, 13, 22, 27, 65–7, 71–4, 77–9, 82–3, 87n22, 90, 94–5, 96n2, 98, 101, 117, 120n5, 122–4, 132, 135, 137–8, 161, 183n4
- Avogadro's 50
- countable 87
- decimal 21, 84, 87n22, 120n4
- encoding 21
- finite 18–19, 24, 64, 69, 84, 117, 161
- hyperinteger 41–2
- hypernatural 18, 42
- hyperreal 41–2, 65, 192n15
- infinite 18, 42, 65
- infinitesimal 41–2, 64–5, 72–3, 76–7, 130. *See also* probability, infinitesimal
- integers 5, 12–13, 16–17, 41, 65–6, 70, 71n8, 79, 83–6, 111, 117–18, 131, 173, 182n1
- Law of Large 80, 170n2
- natural 5–6, 13, 15, 17–18, 20–3, 64, 71, 84, 86, 95, 101n5, 106, 120n5, 123, 129, 164, 183
- prime 12–17, 20, 129n10, 150, 164
- rational 87n22
- real 5–6, 41–2, 64–5, 71n8, 76, 78, 86–7, 101n5, 114, 119–22, 132, 172, 179n7
- surreal 65
- theory of 20, 22
- omnitemporality 171
- organic wholes 51–2, 165
 - infinite 165
- overdetermination 28, 33–7
- paradox 1–37 (*passim*), 140, 142–3, 147, 150–4, 159–60, 166, 180, 189–90, 193
 - Axiom of Choice paradoxes 117–39
 - Banach-Tarski 118, 121, 125, 127, 129, 134, 136
 - Beam's 93, 111–15
 - Benardete's Boards 59, 146–7, 165–6
 - coin-guessing 129–30, 163
 - conservative answers to 1–2
 - decision-theoretic 3, 93–116, 137
 - Gabay-O'Connor 97, 120, 125, 133–4, 138
 - Grandfather 40, 61–3
 - Grim Reaper 40, 46–56, 61–3, 96, 140, 143–7, 149–50, 158, 160, 165–6, 183–4, 191
 - in infinite Newtonian universes 56–60
 - logically revisionary answers to 1–2
 - lottery 64–92
 - metaphysical answers to 1–2
 - Mischievous Toggler 62–3
 - Reversed Grim Reaper 48–9, 56
 - Russell 24n23
 - St. Petersburg 68–71

- Satan's Apple 93, 106–15
- Smullyan's rod 58–60, 157–8
- symmetry 68–70
- Thomson's Lamp 1–2, 4, 7–8, 10, 26, 40–7, 55–6, 63, 142, 151, 153–4, 156, 158, 193
- Zeno's paradoxes 2–3, 10, 59, 160
- Parmenides 10n6
- particles 29, 43, 47n8, 50, 54, 59, 86, 90, 125, 154–6, 161, 175–6, 188
 - antiparticles 54
 - fundamental 156
 - point 147, 155
 - θ particles 42–3
 - ψ -particles 122, 124, 130, 137
- Pascal, Blaise 70n6
- Peano Axioms 18, 20–2
- physics 3–4, 130–1, 152, 158, 161, 167, 172–8, 180, 194
 - deterministic 175
 - fundamental 154–5
 - Newtonian 32, 40, 42, 56–60, 63, 90, 154, 156, 172, 174, 176, 178n6
 - particle 155, 176
 - philosophy of 180
 - Quantum Mechanics 90, 172, 174–6
- Plato 29n2
- pluralities 6–7, 19, 29n3, 30–1, 34, 52, 154, 162, 182–3
 - finite 18
 - infinite 13
- possibility 16, 48, 50, 52–3, 57, 60, 71–2, 81, 87, 92, 164, 171, 184, 193
 - causal 123
 - distinct from narrow logical coherence 53
 - metaphysical 3, 7, 17, 23, 37, 40, 42, 61n20, 89–91, 122, 124–5, 131, 137, 152, 155–6, 170, 191–2
 - nomic 89–90, 176
 - of an actual infinite 13
 - of causal loops 63, 194. *See also* causation, causal loops
 - of future infinities 17, 103, 150, 170. *See also* infinity, future
 - of God 181. *See also* God, theism
 - of immaterial beings 154, 158–9. *See also* God, minds, immaterial
 - of infinite causal histories 37, 193
 - of an infinite multiverse 118, 148. *See also* multiverse
 - of infinite regresses 30
 - of Newtonian physics 60. *See also* physics, Newtonian
 - of paradox 103. *See also* paradox
 - of paradoxical lotteries 92. *See also* paradox, lottery
 - of past infinities 151, 170. *See also* infinity, past
 - of time travel 61, 63, 152. *See also* time, travel
 - of traversing an infinite series 152
 - of vicious regresses 29. *See also* regress, vicious
- operator 123
- physical 135
- powerset 5, 15
- presentism. *See* time, presentist theory of
- Principle of Sufficient Reason (PSR) 2, 29, 43–4, 46–7, 137, 164n10, 182, 193
- privations 149, 166
- probability 64–5, 67, 69–72, 74–8, 80, 82, 86, 87n23, 89–91, 102, 104–5, 121, 130, 170, 175–6, 188
 - Bayesian 71–2, 74–5, 92
 - classical 75, 77–8, 80–2, 87n24, 96–7, 138
 - conditional 72, 77
 - epistemic 91, 105
 - function 75–6, 85–6
 - infinitesimal 65, 72–4, 76, 78–9, 82, 85–6, 87n22, 91–2, 96n2, 102, 110, 129
 - measure 81, 138
 - non-classical 82, 129n11
 - posterior 72
 - prior 72
 - rational 76
 - theory 75, 78, 81, 123, 129n11
 - unconditional 124
 - zero 65, 69, 71n8, 72–4, 76, 77–81, 85–6, 87n22, 91, 96–7, 124, 129–30, 137–8
- processes 80, 86–9, 103, 105, 124, 171n3
 - causal 54, 81–2, 87–8, 92, 107, 125
 - continuous distribution 92
 - deterministic 17n15, 40, 193
 - indeterministic 176
 - infinitely complex 89
 - metabolic 142
 - physical 123, 131
 - radioactive decay 87, 169–70
 - random 34, 49, 92, 135
 - stochastic 81–5, 118, 123
 - symmetric 81–5
- proofs 14, 16, 20, 22
 - finite 22
 - infinite 22
 - provability 16–17, 20–2
- properties 9, 11, 17, 21–2, 36, 41–2, 61, 76, 81, 86, 90n25, 95, 98, 122, 129–30, 140–1, 145, 161, 164, 171, 179, 182, 187
 - accidental 187
 - Archimedean 41, 64n1
 - distributional 179
 - functional 132
 - intrinsic 9, 187
 - microphysical 141
 - of spatiality 131. *See also* space

properties (*Cont.*)

- statistical 175
- subdivisibility 183

quantification 123

- plural 6

Quantum Mechanics 90, 172, 174–6

radioactive decay 87, 169–70

- rearrangement principles 7, 9, 50, 57, 126
- unrestricted 8, 50

reductio ad absurdum 5n4, 13, 18, 51, 87n23, 98, 167, 170, 182

Reflection Principle 76

regress 62, 140, 142, 163

- backwards 26, 31
- Bradley Regress 29
- causal 25–37 (*passim*), 151, 162, 193
- dependence 27–8
- explanatory 29
- grounding 33, 36
- infinite 2–3, 25–39, 62–3, 92, 151
- justification 27–8, 34
- non-vicious 27–8, 33, 35–6
- set membership 36, 162n8
- Third Human Being Regress 29
- uncaused 27–32, 35
- vicious 27–37 (*passim*), 187

relations 18–19, 38–9, 143, 163, 181, 190

- antisymmetric 19
- causal. *See* causation, causal relations
- counterfactual 54, 190
- equivalence 84, 98, 102, 120, 134–5, 139
- explanatory 164–5, 187, 190
- justification. *See* justification, epistemic
- reflexive 98
- total 19
- transitive 19, 142, 146

Relativity Theory 150, 155, 167

- General 157, 169, 178, 180
- Special 42–3, 47n8

Riemann Zeta Hypothesis (RZH) 17n13

Russell, Bertrand. *See* paradox, Russell

sequences 15, 21, 32–3, 44–5, 50, 63, 67, 79–84, 88–9, 113–14, 123–4, 137, 143, 160, 163, 168

- causal 4, 7, 32, 37, 49, 87, 147–8, 182, 194
- countable 45
- dense 32, 37
- finite 21, 37–8, 162, 186
- infinite 10, 20, 27, 32–3, 38–9, 45–6, 49, 56, 80, 82–7, 94–8, 115, 120–2, 129–30, 132–4, 137–8, 143, 147–8, 152, 154, 183–4
- monotonic 162–3
- theory of 174
- time 61, 122, 169–72

series

- convergent 76n13, 111–12
- Laurent 65, 76n12

set 4–6, 12, 15, 18, 20, 23–4, 36, 38, 41, 77, 79, 81–4, 88, 96–8, 100, 102, 107, 112, 115, 117–24, 127–9, 133–6, 138–9, 162n8, 179n7

- cardinality 5–6, 18, 23–4, 100–1, 119–22, 132, 139

- countable 3, 6, 18–23, 50, 87, 101, 119, 121–2, 124, 132, 138, 193

Dedekind-finite 18n17

Dedekind-infinite 5n3

empty 5, 20

- finite 5–6, 12, 18–23, 38–9, 79, 101, 112, 115, 117, 171, 193

- infinite 5–6, 12, 15, 101n4

- measurable 121, 123, 136

- non-measurable 82, 120–1, 125, 134, 136–7

- powerset 5, 15

Russell 24n23

- uncountable 6, 15, 42, 101

set theory 3, 14, 18, 20, 22, 25, 36, 83, 91, 117–18, 136

Zermelo-Fraenkel axioms of. *See*

Zermelo-Fraenkel axioms

space 57, 125–6, 129, 131, 147, 155, 160–1, 168–9, 179, 183

- as relational 57n13

- continuous 131–2, 160, 167, 172, 178, 180

- discrete 4, 87, 147, 160, 167–9, 172, 178–80, 194

- empty 158, 179–80

- Euclidean 121, 179n7

- finite 160

- four-dimensional 131–2

- functional view of 131n13

- Hilbert 89–90

- infinite 86

- Newtonian 60. *See also* physics, Newtonian

- points of 169

- sample (in probability) 74, 76, 81–3, 115, 126–7, 129, 138

- vector 90n25

spacetime 7, 9, 23, 63, 131, 140, 144, 159–60, 170, 172. *See also* time

- curved 169

- eleven-dimensional 131

- five-dimensional 132

- flat 178

supertasks 40–63, 66, 71n8, 79–81, 83, 85, 87n23, 88, 108, 112–14, 131, 134–5, 151, 159

- reversed 85

symmetry 66–71, 77, 81–5, 92, 104–5, 118

- epistemic 105

- ontological 175

- testimony 27–8
 - divine 135
- theism 4, 8, 181–2, 184–92, 194. *See also* God
 - classical 8, 152, 181, 185, 187–8, 192
 - open 186
- theorem 25, 38, 76n13, 77, 95, 102, 115, 118, 136n17, 138
 - Banach-Tarski. *See* paradox, Banach-Tarski
 - Bayes' 72, 74–5
 - Cantor's 5, 15
 - Fermat's Last 20
 - Gödel's incompleteness theorems 16, 18, 20–2, 118n2
 - Good's 75
 - Isaac Newton's shell 57n14
 - Schröder-Bernstein 5, 120n3
 - Soundness 20
- time 1–2, 10–11, 17, 23–4, 29, 41–52, 54–5, 61–2, 85, 87, 103–6, 110, 113, 122, 131–3, 135–6, 143–4, 146, 150–3, 156, 159, 161, 167–80, 182n2, 183–4, 191
 - as the measure of change 168, 171
 - beginning of 122
 - collapse 178. *See also* physics, Quantum Mechanics
 - continuous 26, 87, 131–2, 160, 167, 167–80 (*passim*)
 - derivatives with respect to 178n6
 - direction of 146
 - discrete 1, 4, 46, 55–6, 87, 152, 160, 167–80, 194
 - eternalist theory of 10–11, 17, 23, 49–50, 103, 150, 193
 - external 60, 152, 170–2
 - fictitious 178
 - four-dimensionalist theory of 164
 - growing block theory of 3, 10–11, 17–18, 23, 150, 193
 - infinite 55, 85, 103, 112, 122, 159, 192
 - infinitesimal 169, 170n2
 - instants of 43, 47–8, 156, 159–60, 167–8, 170–1, 177–8, 183
 - internal 60, 170–2
 - lifetimes 113, 124, 130, 137
 - paradoxical 54–5
 - Planck 174
 - possible 87, 183
 - presentist theory of 10–11, 17, 23–4, 150
 - spacetime. *See* spacetime
 - timelessness 54
 - travel 3, 40, 61–3, 152, 162–3, 170
- tropes
 - abundant 141
 - individuation of 140–2
- truth 16–17, 98, 164n10
 - contingent 187
 - future 186
 - mathematical 16, 20
 - necessary 104–6, 161
 - particular 164
 - universal 164
 - unprovable 16. *See also* theorem, Gödel's incompleteness theorems
- universe 8–9, 23, 26–7, 33, 37, 49, 132, 135, 137, 153–4, 161, 167, 173, 175, 180, 188
 - causally isolated island 50, 122–4, 130, 148, 152
 - cause of 54. *See also* God
 - final state of 146
 - finite 91, 177
 - Grothendieck 136n16
 - infinite 153, 177
 - infinite Newtonian 56–60
 - quantum 177
- utilities 109–11, 113–14, 154
 - expected 68–71
 - finite 69–70
 - hyperreal 192n15
 - infinite 154
 - maximization 69–70
 - subjective 154
- van Fraassen, Bas 76
- van Inwagen, Peter 52
- vectors 90, 154, 156
 - vector space 90n25
- Young's modulus 157–8
- Zeno of Elea. *See* paradox, Zeno's paradoxes
- Zermelo-Fraenkel axioms 20, 36, 118–19, 136